

On LFY Discs: openSUSE 11.2, Mandriva 2010 & Ubuntu 9.10

Rs 110
ISSN 0974-1054

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 07 ISSUE: 10 December 2009 116 PAGES ISSUE# 83

ForYou

An **EFY**GROUP Publication
ISO 9001:2008 CERTIFIED



NetBooks

Facts, Figures, Options &
Opinion

Other features inside

**Private Cloud Computing
with Ubuntu Enterprise Cloud**

**Develop Games using Open Source
Graphics Engines, Libraries & More**

**Improve Performance on Multicores
By Adding Threads to Existing Code**

**Counter Disk Fragmentation
& Improve RAID Performance**



(K)Ubuntu 9.10
openSUSE 11.2
Mandriva 2010

Peek inside, for
**the Ultimate
Distro
Showdown**

India INR 100
US \$ 12
Singapore S\$ 9.5
Malaysia MYR 19





Stop SPAM

Stop Virus

and...

Stop Paying!

Best
Welcome to the World's ~~First~~ Free Email Security Service

To activate SAFENTRIX for your domain, Visit <http://www.safentrix.com> now!

Tyrone®

BUILD YOUR APPLIANCE



You might have a perfect idea, or the software for that matter but you might still struggle for a perfect hardware to sell your idea in an appliance form.

We offer a wide range of flexible hardware platform which is reliable, powerful and cost effective. Please do get in touch with us with your requirement.

Netweb
TECHNOLOGIES

ISO 9001:2000 Company

Corporate Office
B-1/A25, 2nd Floor
Mohan Co-Operative Industrial Estate
Mathura Road, New Delhi - 110044

New Delhi : +91 11 29942258/29942297
Bangalore : +91 80 41146565/40540000
Chennai : +91 44 43531009/26631975
Kolkatta : +91 33 23232055/23232056
Singapore : +65 64157239/92288638

Website : www.netwebindia.com
www.tyronesystems.com

Email : info@netwebindia.com
netweb@vsnl.com



34

NetBooks

Facts, Figures, Options & Opinion

Peek inside for
The Ultimate
Distro
Showdown

(K)Ubuntu 9.10
openSUSE 11.2
Mandriva 2010

20

FOR YOU & ME

- 20 The Ultimate Distro Showdown
- 30 FOSS in Action: Preparing for that Crucial Meeting
- 34 NetBooks: Facts, Figures, Options & Opinion
- 67 It Plays Music, At Your Command (Line)
- 70 Explore the Save Mode With FOSS
- 84 Up, Close, & Personal with the creator of OLPC XO eBookReader



DEVELOPERS

- 76 Think Parallel with OpenMP
- 80 Getting Started with Game Development
- 92 Create Cute Apps with Qt, Part 2
- 98 SubConf 2009—A Report
- 100 Subversion Version Compatibility



CONTENTS

ADMIN

52 Private Cloud Computing with UEC

62 Performance Tuning and Monitoring—Part1



GEEKS

73 Getting Started with GNU Octave—Part1

COLUMNS

83 The Joy of Programming: Parsing Complex Declarations in C

88 Python in Research, Part 4—Optimising Returns, Minimising Risk

96 CodeSport

102 A Voyage to the Kernel—Part 19:Day 18: Segment 3.8

REGULAR FEATURES

06 Editorial

08 You Said It...

10 Q&A Section

12 FOSS Bytes

50 CodeChef

60 Tips & Tricks

106 Linux Jobs

108 FOSS Yellow Pages



LFY DVD: openSUSE 11.2 & Mandriva 2010

"Killer looks, slick integration and superb tools makes it THE choice for all but novices or very finicky Linux users out there"

"The large amount of software included on the DVD means you don't need an Internet connection for anything except updates"



LFY CD: Ubuntu 9.10

"Perfect for people who want DEBs at all costs and want the system customised exactly their way"

Turn to Page 20 to learn how the three distros fared when pitted against each other.

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgroove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhhd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 022- 40497401, 40497402,
40497474, 40497479, Fax: 40497434
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear Readers,

Will Ubuntu become history now?

That's the question that popped in my mind when I read through the Google official blog on November 19—the day it open sourced its Chrome OS. Of course, the official title given to the open sourced version is Chromium OS—I guess Google has learnt a few things from Red Hat, when it comes to open sourcing projects.

The Chromium OS project includes the "...current code base, user interface experiments and some initial designs for ongoing development," according to Google. Why go open source now? To connect with the open source community and get it involved, obviously. What else did you expect?

To me, the Chrome OS can usher in a new era—one that (hopefully) will take Linux, open source and cloud computing to the masses. If this happens, many users, especially the new ones, may never recognise that shift, since this will then be the *de facto* way of life for them, just like Windows is today, for many.

Chrome will surely accelerate the demise of traditional client-server applications, and all that an application will need to interact with is the Web browser. Of course, all those who think that this will make life simpler for developers (as they won't have to worry about the underlying hardware)—think again. I have met many developers who took great pains to develop Web-based Taj Mahals, and then despaired when they realised that their creation looked dilapidated on other browsers.

So, how good is the Chrome OS? The answer will depend on whom you ask. Fans of Google will swear by it. While others simply consider it 'useless' right now. But then, the final version of this OS is expected only in the second half of next year (2010). So, if you can find flaws and report them, you would be helping the core team. It would be better still if you could contribute in developing some of the features on the 'to do' list.

Is Chrome an upgrade of the Android? No.

Android was designed from the beginning to work across a variety of devices—from phones and set-top boxes to netbooks. The Chrome OS is being designed to power computers ranging from small netbooks to full-sized desktops. While there are areas where the two overlap, they will be two separate open source projects powered by Google.

Back to our first question—what happens to Linux distros like Ubuntu that have gained popularity amongst end users? Will Chrome displace them? I hate to make wild predictions, but the gut feeling is that Google has most often got things right—whether it was with its search engine, with Google Docs or even Gmail. Plus, no one can refute the increased popularity of Web-based applications.

Chrome will surely accelerate the demise of traditional client-server applications, and all that an application will need to interact with is the Web browser.

Whatever be the outcome—two things are certain. Open source will emerge as a winner, and we will be there to report that.

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 38 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Qmail Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration
LLC410: Nagios - System & Network Monitoring Software
LLC412: Apache & Secure Web Server Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC209: Joomla CMS
LLC501: Programming with OpenGL
LLC504: Linux on Embedded Systems

RHCE Certification Training

RH033: Red hat Linux Essentials
RH133: Red Hat Linux System Administration
RH253: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCSS / RHCA Certification Training

RHS333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RH401: Red Hat Enterprise Deployment & Systems Management
RH436: Red Hat Enterprise Clustering & Storage Management
RH442: Red Hat Enterprise System Monitoring & Performance Tuning
RHS429: Red Hat Enterprise SELinux Policy Administration

RHCSS / RHCA Courses & Exams

RHS333: 5 Dec.; RH423: 19 Dec.
RH429: 19 Dec.; RH401: 1 Dec.
RH436: 7 Dec; RH442: 14 Dec.
EX333, EX423 on 18 & 22 December
EX401, EX436 & EX442 on
5, 11 & 18 December respectively

Joomla CMS Course on 19 & 20 December 2009

RH301 from 7, 14, 21 & 28 Dec. '09
RHCE Exam in Bangalore on 22 Dec. '09

**LLC504: Linux on Embedded Systems
4 Day Fast Track Course
starting on 26 December 2009**

**LLC410: Nagios System & Network
Monitoring Software - Training from 12 Dec '09**

Linux Support & Solutions

Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS Servers

For more info log on to:

www.linuxlearningcentre.com

Call: 9845057731 / 9343780054

Email: info@linuxlearningcentre.com

RED HAT Training Partner
RHCE, RHCSS & RHCA Exam Centre



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

635, 6th Main Road, (Adj.. Bank of India) Hanumanthnagar, Bangalore 560019

Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839

Cell: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com

www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore

1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070

Tel: +91.80.26712928

You said it . . .



I recently read the article 'How to Crack Job Interviews' in the October issue of LFY. While the article addresses the subject admirably, we need to also address the question of skill deficiency. It is a well-known fact that only about 10 per cent of job aspirants get through the interviews. The selected few generally go through a training program, which polishes their IT skills and improves their soft skills. The main focus of these training programmes is 'learning by doing'. The other 90 per cent do not make it because they lack IT skills, and many lack soft skills too. There is a need to address the issue of the inadequate IT skills of our graduates. Linux can play an important role in this area.

With hardware prices falling, a laptop or a desktop is within the reach of many students. Install any standard Linux distribution, and you get a full-fledged system with many compilers, libraries, tools, etc. From then onwards, depending on an individual's interest, any type of programming can be learnt. In fact, the choice is so huge, it can be very overwhelming! If one is interested in traditional languages, C and C++ are available. If one is interested in functional programming, 'SCHEME', 'HASKELL', and 'ERLANG' are available. If your interest is in Web development, the usual LAMP stack is available. However, for programmers who are beginners, I would recommend Python as a starting language, since it has a simple syntax and excellent support in terms of libraries.

Often, I have heard students saying that they don't have problems to solve, which is unfortunate, because there are many good repositories freely available on the Internet. Here are a few sites I know of:

- projecteuler.net
- www.inf.bme.hu/contests/tasks
- www.karrels.org/Ed/ACM
- uva.onlinejudge.org

So, in conclusion, I would like to see Linux being used, more and more, as a learning tool, in addition to being used for many other purposes.

—**Lalitha K Prasad [lalitha.prasad@tcs.com], head, Corporate Learning Centre, Tata Consultancy Services**



I am an associate professor in medicine at a medical college. I discovered Linux two-three years back and now try to follow and support OSS/Linux in every way possible. I have converted around 12 doctors into Linux users and fans. I had been seeking some advice on an Android phone since October last year. Please accept my heartiest congratulations on your story, 'Android Experience' (in the November 2009 issue). However, I do feel that comparing HTC Tattoo (www.htc.com/www/product/tattoo/overview.html) with HTC Hero would have served readers better. HTC Tattoo would have satisfied more readers and buyers than HTC Hero, with its bloated price tag. We now await reviews on affordable Android phones!

—**Dr Smit Shrivastava, dr.smit.shrivastava@gmail.com**

The author, Sayantan Pal, replies:

First and foremost, I would like to express my elation for your kind words. About the comparison between HTC Hero and HTC Tattoo, the Hero was launched in the first week of September 2009, in India, while HTC Tattoo was launched in India on November 18 (despite being launched in European countries in the first week of September itself). It was because of the timing of the

Indian launch that we couldn't come up with a comparative review.



I am a regular reader of *LINUX For You* and am very impressed by its contents and presentation. Nagios seems to be a good system and network monitoring software. I have tried it many times but I do not know how to use it effectively. Could you please help systems administrators like me by publishing an article on how to install and use Nagios?

—**S Chinnachamy, chinnachamy.s@gmail.com**

ED: We've covered Nagios in depth back in December 2007... Hope that article (sent to you as a PDF) has helped you to get started.



I am eagerly awaiting the December issue of *LINUX For You* hoping it will come bundled with Mandriva 2010. Because of low bandwidth, it is very difficult to download a 4 GB operating system. The best way to get it is through your magazine; so please include Mandriva 2010 in your next issue.

—**Niranjana Adiga, nvadiga@yahoo.co.in**

ED: Well, as you can see, the December 2010 issue has come with Mandriva 2010. And what's more, there's Ubuntu 9.10 and openSUSE 11.2 as well :-)

Please send your comments or suggestions to:

The Editor

LINUX FOR YOU Magazine
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: 011-26810601/02/03
Fax: 011-26817563
Email: lfyedit@efyindia.com
Website: www.openITis.com



ElectronicsForYou

presents

EFY SUMMIT | 2009

'Unleashing New Opportunities: Indian Electronics Industry'

December 16, 2009 • Hotel Le Meridien • New Delhi

EFY SUMMIT 2009: A FIRST-OF-ITS-KIND EVENT THAT WILL BRING TOGETHER LEADERS FROM THE INDIAN ELECTRONICS INDUSTRY UNDER ONE ROOF

All industry reports state loud and clear that there's tremendous opportunity for growth in the Indian electronics industry. The **EFY Summit** will help decision-makers to **unleash new opportunities** and devise strategies to tap them.

The Agenda

- Which products or customer segments promise growth
- Which markets remain untapped
- What's new in terms of technology for you to harness
- What is it in the Indian eco-system that you can benefit from
- How to fund your expansion
- When to go for an IPO

Why You Should Attend

- Discover 'where' and 'how' to expand your business
- Discover 'untapped' markets and product segments
- Identify technologies that can change your industry
- Identify technology trends that can affect your markets
- Learn from the success stories of others
- Meet and network with top-shots of the electronics industry

Who Should Attend

- Heads of organisations (CEOs, Managing Directors, Partners, and others)
- Senior heads of business units (Vice Presidents, General Managers, Heads of R&D Units, and others)
- Senior sales and marketing professionals
- Policy makers for the ICT and electronics sectors
- Infrastructure solution providers for the ICT and electronics sectors
- Representatives of banking and finance companies
- Anyone with an interest in the fast growing electronics industry

Book Your Seat Now!

Delegate Fee: Rs 5000* plus 10.3% tax; Early Bird Discount: 10% (till 15th Nov, 09)
To register, call: 09312888878 OR email at: efysummit@efyindia.com

** Special group discounts available*

For Sponsorships and Other Branding Opportunities

Call 09312888878 OR email at: efyenq@efyindia.com

Media Partners

ELECTRONICS
BAZAAR

Electronics Bazaar

EFYtimes.com
IF ITS TECHNOLOGY ITS HERE

EFYTimes.com

ElectronicsForYou
Annual Guide

Electronics For You
Annual Guide

eleB2B.com

Eleb2b.com

Supported By

ELENA

CMAI

MAIT



KNOW HOW



Q I am new to Linux and have recently installed it on my computer. I want to know if I can run Windows programs and games on it? If yes, then please explain how to go about it.

—Nordon Lama, Sikkim

There are emulators that can run some of the Windows programs and games on Linux. WINE is perhaps the best amongst the lot. To get started with the installation, the best option is to download the precompiled binary package for your Linux distribution from www.winehq.org/site/download. Note that not all Windows-based apps will run.

Q I have been using Thunderbird 2.0 as my main e-mail client. I have set about 35-40 message

filters for my incoming mails. A few days back I bought a new laptop and moved all my data, including the mails. Now the problem is that I am unable to export my filters to the new set-up. Is there any option in Thunderbird that allows me to export/import message filters or do I have to go through the pain of creating each of them again?

Thunderbird does not have any function built-in to import or export message filters. There is a Thunderbird extension available [<https://addons.mozilla.org/en-US/thunderbird/addon/2474>] that has the functionality you want. You need to download the .xpi file and install the extension on Thunderbird, on both your systems.

After installing the extension, open *Tools→Message Filters*. Click *Export Filters* to export the message filters for the selected account. Copy the file over to the new system and again go to *Tools→Import Filters* to import message filters.

Q I am facing a problem with OpenOffice.org. Whenever I create a new document, it saves it in the .odt format, by default. This creates problems for other people who use Microsoft Office. Can you help me in setting my OpenOffice.org to

save documents by default in the .doc format?

—Sneha Basu, Kolkata


To set OpenOffice.org to save documents in the .doc format by default, follow the steps given below:

1. Go to *Tools→Options* to open the *Options* window.
2. Here, on the left pane, expand the category called *Load/Save*, and click on the *General* menu.
3. Now look for an option called 'Always save as' and select your option there.

From now on, your new documents will be saved in the .doc format by default.

Q I am using Firefox 3.0 and whenever I type a few characters of a URL, Firefox suggests a URL from the list of previously visited sites in the address bar's drop-down history menu. I want to remove some URLs from the list of previously visited sites. Is it possible to do so without deleting all?

—Parkash Singh, Ludhiana

You can just use the down arrow key to select the URL you want to delete—it'll appear in your location bar. Instead of hitting *Enter* to go to the website, use the *Delete* key. Firefox won't suggest the same URL again. **END** 

More and more IT Managers are adopting Open Source. What about You?

Read LINUX For You: Asia's Leading Open Source Magazine

www.linuxforu.com



Want to be the Best ETHICAL HACKER

&

Is Your Computer Secure? View Your Security Through the Eyes of a Hacker



The Secret of Hacking :: 2nd Edition

Over
50,000
Sold!



First Edition FREE (PDF Version)

After the grand success of the first edition that came out in June 2009 Leo Impact has come back with a 4 times more powerful second edition.

Even the most secure computers are Hackable...

- All E-mail addresses are Hackable, including Gmail, Yahoo!, Rediff etc.
- All PCs can be hacked remotely using the latest tools and exploits.
- Even mobile numbers can be spoofed (Show any Mobile No on Victim mobile)
- Easily pass CEH (ver 6), CHFI, CISSP, CISA Certification.
- Learn how to secure your system and network from hackers.
- Learn Advanced Ethical Hacking:
- Metasploit & Backtrack & Untraceable Hacking
- Advanced Penetration Testing & Vulnerability assesment.

When you buy the book, you get:

- 1 Printed Book (Second Edition) + First Edition (PDF)
- 2 DVD (18,500 tools, e-books, videos)
- E-mail Technical Support
- Free Lifetime Membership to Access Videos & Tools

Why us ?

- * We provide email support and free Lifetime Membership.
- * Each Topic cover by Video Demonstration.
- * 30 days money back guarantee and Free Shipping Cost.



Payment modes:

Credit Card, Netbanking, Cash Deposit & DD/Cheque

To Order Our Book Just SMS **Leoimpact Yourname to 54242**

For more info. & online order: www.theseretofhacking.com

Help? Call us now !! +91.9829944518, 0141-4039187, +91.9953244518, (10:30 AM to 11 PM)




LEO IMPACT SECURITY SERVICES PVT LTD

T8, Malyia apartment, near BJP office
c-schme, jaipur (Rajasthan) 302001
Email: contact@leoimpact.com

2029 Century Park East, 14th Floor,
California 90067 United States
Phone : +1 818 252 9090(USA)

We welcome Partners, Investors, and Distributors :: We conduct Hacking workshop Also.



Linux to overtake Windows on netbooks by 2013

New research projects that in 2009 Linux will represent 32 per cent of global netbook sales, far higher than the seven per cent figure claimed by Microsoft. The market data from ABI Research estimates that Linux will overtake Windows on netbooks by 2013, largely due to sales in less-developed countries, as per DesktopLinux.com.

ABI Research's study, *Netbooks, MIDs and Mobile CE Market Data*, says 35 million netbooks will be shipped by manufacturers in 2009. Asus dominated the netbook category in 2007 when it debuted the Eee PC, and Acer made a big push in Q4 2008 to lead the second year.

"74 per cent of 2008 netbook shipments bore the brand names of just three vendors—Acer, Asus and Samsung," explained Jeff Orr, senior analyst, ABI. "However, the rapid growth of netbooks as a second computer in developed markets will be eclipsed in the coming years by vendors targeting developing nations offering netbooks as the first Internet PC at homes."

ABI Research forecasts UMDs (ultra-mobile devices --a combination of netbooks, MIDs and UMPCs) to reach 124 million systems in 2011.

Orr believes laptop vendors may soon "reinvent the UMPC," bringing down the low-end laptop price to \$500 through ultra-low voltage machines. "That could affect some netbook sales in developed markets, especially among the business users."

As netbook markets thrive, the mobile consumer electronics category, which is expected to account for only about two million device shipments this year, is forecast to top 50 million in 2014. This market is currently led by connected personal navigation devices (PNDs) and eBook readers. Amazon's Kindle has lots of competition coming from the Barnes & Noble Nook, Sony Daily Edition, iRex, Plastic Logic Que and others.



Adobe AIR 2 beta for Linux, released

Adobe Systems Incorporated has announced that beta versions of Adobe AIR 2 and Adobe Flash Player 10.1 are now available. According to the release information, with AIR 2, developers can create more reliable and feature-rich applications in the desktop context, and they can easily add AIR 2 functionality to update current AIR 1.0 applications.

New AIR 2 features include enhanced support for mass storage devices and native application processes, as well as peer-to-peer and UDP networking. Flash Player 10.1, on the other hand, is supposed to be the first consistent browser runtime release of the Open Screen Project that will enable uncompromised Web browsing of expressive applications, content and high definition (HD) videos across the desktop and devices. The beta versions are available as free downloads from Adobe Labs.

A racing bike with Linux inside

Mavizen, the manufacturing arm of the TTXGP, recently announced the TTX02 at this year's SEMA. The Mavizen TTX02 electric motorbike is shipping not only with a chassis and drivetrain, but with an open source Linux OS, Web server, USB-based system bus and Wi-Fi connectivity. Production will be limited to just 50 units when the Mavizen TTX02 arrives sometime in the first quarter of next year, reports Engadget.

White House goes Drupal

WhiteHouse.gov, the official website for the White House and President Barack Obama, now runs on open source. The White House has launched a new version of its website, powered by the Drupal content management system and the Apache Solr search server.

"We wanted to improve the tools used by thousands of people who come to WhiteHouse.gov to engage with White House officials, and each other, in meaningful ways," White House spokesman Nick Shapiro said in an e-mail to InformationWeek. "We now have a technology platform to get more and more voices on the site. This is state-of-the-art technology, and the government is a participant in it."



Searching for the world's most secure and fastest web hosting services



Everything you need for Internet success:

•Domain Registration & Web Hosting •Dedicated Servers •Reseller Hosting

Lifetime Plan

- 1 Free Domain Lifetime*
- Unlimited* Web Space + Cpanel11
- Free Account Transfer & Migration
- Unlimited E-mails & FTP Accounts
- PHP 5, ROR, Perl, Cron, Real Statics
- Unlimited Mysql & Postgree Database
- Fantastico DeLuxe (52 Scripts)

Pay **Rs. 5,000/-** for lifetime

Reseller Hosting

- 500 GB Dedicated Space with Overselling Feature
- Unlimited Space, Bandwidth, E-mail A/C, Database,
- You can create Unlimited CPanel Hosting Account
- Free Billing Software with Integration
- Free Paypal account to accept Credit card Payments
- Free Hosting Account Transfer from Old Hosting
- Free Private Name Server (ns1.yourdomain.com)

Pay **Rs. 5,000/-** for 6 months

Web Hosting (CPANEL)

- 1 Free Domain for Lifetime*
- Unlimited* Web Space + Cpanel11
- Free Account Transfer & Migration
- Unlimited E-mails & FTP Accounts
- PHP 5, ROR, Perl, Cron, Real Statics
- Unlimited Mysql & Postgree Database
- Fantastico DeLuxe (52 Scripts)

Pay **Rs. 2,150/-** for a year

Note: Pay only Rs. 500/- as Domain Fee and use our hosting services free of cost for 7 days.

Other Services

- Website Designing
- Bulk SMS and E-mail marketing
- Computer Security Services

What Make Us Different from others

- INSTANT Activation + Perfect balance between great prices and high quality services
- We are available 24x7x365 on live chat & web call system + Free basic management
- 100% network uptime & 99.9% server uptime with service level agreement
- SSR Technology (Highest speed, security and uptime on Internet)

Payment modes:

1. Credit Cards
 2. Demand Draft / Cheque
 3. Cash payment in ICICI
- www.kghosting.com/payments



For more info. & demo Just: **SMS Leoimpact Yourname to 54242** Or log on to: **www.kghosting.com**
Help? Call us now !! **+91.9829944518, 0141-4039187, +91.9953244518** (10:30 AM to 11 PM)



**LEO IMPACT
SECURITY**


LEO IMPACT SECURITY SERVICES PVT LTD

Corporate Office:

T8, Malyia apartment, near BJP office
c-schme, jaipur (Rajasthan) 302001
Email: contact@leoimpact.com

USA :

2029 Century Park East, 14th Floor,
California 90067 United States
Phone : +1 818 252 9090(USA)



Google previews Chrome OS, releases source code

Google gave a preview of its upcoming Chrome operating system in the US on November 19. Scheduled for launch in a year, the open source OS is intended for use on low-cost netbooks, but the final version will only be available pre-installed on new devices. Google has a number of prerequisites for hardware manufacturers to ensure Chrome's smooth running.

In essence, Chrome OS is a much-expanded version of Google's Chromium Web browser and almost all of its functionality is based around Web applications. Although the OS code is stored on the netbook itself, the applications it runs are all cloud-based—Gmail, Google Docs, and so on.



Not surprisingly, Google Chrome OS also relies upon a ubiquitous Internet connection, but most Google apps already support offline access via Google Gears. How much of a problem a Chrome OS netbook is when there's no Internet connection remains to be seen, but

it's impossible to imagine that this isn't a situation that Google hasn't already thought about.

Red Hat Virtualisation for heterogeneous servers

Red Hat has announced the availability of Red Hat Enterprise Virtualisation for Servers, the newest product set in the Red Hat Enterprise Virtualisation portfolio. Red Hat Enterprise Virtualisation for Servers is designed to enable pervasive adoption of virtualisation, with an end-to-end solution combining a standalone hypervisor and powerful virtualisation management.

Paul Cormier, executive vice president and president, products and technologies, Red Hat, said, "Red Hat Enterprise Virtualisation provides ISVs, IHVs and customers with the opportunity to seamlessly move applications and hardware platforms from bare metal computing to virtualised computing to cloud computing while maintaining certifications. We believe that Red Hat Enterprise Virtualisation, especially the management components, will allow customers to use virtualisation pervasively."

Red Hat Enterprise Virtualization for Desktops remains in private beta today and is expected to be made generally available in early 2010.

FOSS jobs portal launched

LUG@IITD community manager, Narendra Sisodiya, announced FOSS job board (<http://fossjobs.in>) on the 20th of November, '09. There was a high requirement of a job board for students and professionals who want to work on free and open source technologies. FOSSJobs.in is committed to provide, free of cost, job postings which are strictly related to FOSS. The website has a nice GUI to search jobs by category, by city and by company. You can also search for jobs by the type (freelance, work from home, full time, part time, volunteer). All updates on fossjobs.in are available on Twitter too (@fossjobsdotin).

Dell unveils Android-based Mini 3 smartphone

Dell has confirmed plans to enter the smartphone business. The company officially announced recently that it would sell its first smartphone in China in late November, followed by a launch in Brazil towards the end of the year. The much talked about new phone from Dell is called Mini 3 and uses Android as its OS. However, the PC manufacturer didn't disclose further technical details about the phone. Dell said it would distribute its new Mini 3 smartphones through China Mobile, the largest telecommunications company in the world with more than 500 million customers, and Claro, which serves more than 42 million people in Brazil as part of the America Movil network. The company said it chose Android because of its easy customisation, explaining, "Mobile operators like it because they can easily customise the operating system in order to make it work better with their wireless platforms."



Netweb unveils Tyrone Opslag flexible storage solution

Netweb Technologies—a server, storage and HPC solutions provider—has unveiled its modified flexible storage solution (FS2) for the broadcast and media industry that is expected to significantly change traditional SAN and NAS markets. Key features include scalability up to 384 TB, transfer speeds up to 2GB/s when used with InfiniBand, and input/output unification. FS2 can store more than 900 hours of uncompressed, high-definition 1080i video footage and 30,000 hours of uncompressed standard-definition (SD) video. Besides this, a GUI and Web based management system makes FS2 deployment easy and fast.

Become an IT Expert

Join

USS

Ultramax System Specialist

100 % job Written Guarantee

Other courses & certification

SAS Certification
with exam preparation

Clinical Research
Live Project & Internship

RHCE & RHCSS

MCSE

CCNA

CCNP

Sun Solaris 10



Thane
6766 9228

Pune
4107 1111

Bengaluru
4211 5220


ULTRAMAXTM

INFONET EDUCATION

Dadar 6766 9222

enquiry.dadar@ultramaxit.com

[Www.Ultramaxit.com](http://www.Ultramaxit.com)



Fedora 12—it's Constantine this time

Fedora 12, codenamed Constantine, has been released on November 17. New enhancements include next-generation Ogg Theora video, virtualisation improvements and advancements to NetworkManager, amongst numerous others.

The new Theora 1.1 enhances the quality of free video to better meet user expectations.

Ogg Theora video in Fedora 12 delivers vibrant media, enabling users to stream and download high-definition video while using 100 per cent free and open software, codecs and formats.

NetworkManager improvements in Fedora 12 make it easier

for users to be online using mobile broadband, or configure servers.

Virtualisation-related features include better virtual disk performance and storage discovery, hot changes for virtual network interfaces, reduced memory consumption and a modern network booting infrastructure.

You can get it from <http://get.fedoraproject.org> or wait for LFY's January 2010 issue to get a copy of the DVD.



Project Open 3.4 makes it simpler to implement ITIL

Project Open, a company founded by German engineers in Spain, has released version 3.4, which includes several new packages for the integrated management of IT operations. These new functionalities make it much simpler to implement ITIL and similar best practices, by covering a large percentage of ITIL processes with pre-defined and pre-configured functionality.

All new packages are integrated with project and financial management, making it possible to apply Activity Based Costing (ABC) and other control methods to IT operations. The new version includes: a helpdesk package that supports issue, problem and change management; a configuration database that supports configuration and licence management; a release management and a CVS integration package that support in-house software development, together with the standard PO project and portfolio management.

All packages are based on the PO configuration system, allowing users to add new attributes to objects and configure drop-down boxes. Workflows make it possible to control the life cycles of objects. Integrations are available with Active Directory/LDAP, Nagios, OCS Inventory NG, GanttProject and OpenProj. Integration with SAP and other back-end systems is available on request. Visit www.project-open.org/documentation/install_main for free download and installation instructions.

Firefox turns 5

Mozilla Firefox celebrated its fifth birthday on November 9. The most popular open source browser has achieved many milestones and promises to offer more in the next five years. On November 9, 2004, Mozilla Foundation launched

Firefox 1.0 with the belief that, as the most significant social and technological development of our time, the Internet is a public resource that must remain open and accessible to all.

Within the first four days of Firefox's launch, more than 1 million people had downloaded a brand new browsing experience. In just five years, that number has swelled to over 330 million users worldwide;

almost a quarter of Internet users worldwide use Firefox. Today, Firefox ships in more than 70 languages and offers users more than 7,000 add-ons to help customise their browsing experience.



Event addresses the role of information security

Information security is an integral part of the software development life cycle and needs to be included at every stage. This was the advice conveyed by Securitybyte and OWASP AppSec Asia 2009, an information security event organised in Gurgaon from November 17-20, 2009. Speaking at the sidelines of the conference, Dhruv Soi, chair, OWASP India and partner, Securitybyte, said, "While IT is a dominant sector in India, for millions of developers and IT professionals, information security is still an enigma and area best left to experts." The most interesting talk on open source security was on exploiting Mozilla Firefox extensions. Roberto Suggi Liverani and Nick Freeman, security consultants with security-assessment.com, offered insights into the substantial danger posed by Firefox extensions.



India's ONLY Open Source training institute by
EXPERIENCED System Administrator's & Solution
Architects

**RH 423 with RHDS &
Open LDAP + Active
Directory**

When: 9th to 12th Dec09
Where: FL - Gurgaon
Call: 0124 – 4080880 /
4268187
Mail: info@fosteringlinux.com

**Linux Clusters
With
High Availability**

When: 25th to 27th Dec09
Where: FL - Gurgaon
Call: 0124 – 4080880 /
4268187
Mail: info@fosteringlinux.com

Mail Server

When: 18th to 20th Dec09
Where: FL Gurgaon
Call: 0124 – 4080880 /
4268187
Mail: info@fosteringlinux.com

PHP


When: 16th to 19th Dec09
Where: FL Gurgaon
Call: 0124 – 4080880 /
4268187
Mail: info@fosteringlinux.com

New RHCE & RHCSS
Batches in December
Call: 0124-4080880 / 42681867
Or
Mail: info@fosteringlinux.com

Other trainings at FOSTERing Linux:

- RHCE • RHCSS • RHCA • Jboss
- Clustering • Virtualization • Squid
- Linux Troubleshooting • Mail Server
- Shell Scripting • LAMP • Samba
- PostgreSQL • PHP & My SQL • Perl





Mahindra Satyam joins Symbian Foundation

Mahindra Satyam, part of the \$6.3 billion Mahindra Group, has joined the Symbian Foundation, which, together with others in its ecosystem, is creating an open and complete mobile software platform.

The platform is based on Symbian OS and software assets contributed by Nokia, NTT DoCoMo, and Sony Ericsson, including the S60 and MOAP(S) user interfaces. Portions of the source code are already being moved to open source, under the Eclipse Public License. By mid-2010, this process will be complete. As a member of the Symbian Foundation, Mahindra Satyam gains the immediate right to licence the Symbian Foundation platform, royalty-free and without source code fees; participate in the governance of the foundation; and take part in joint marketing and branding campaigns.

Another Linux-based mobile OS from Samsung

It seems like established Linux-based mobile OSs, Android and LiMo weren't enough. Samsung has announced its plans to come out with its own offering, dubbed bada, in December. The name bada, which means ocean in Korean, was chosen to convey the limitless variety of potential applications that can be created using the new platform, the media release states.

Samsung expects the first handsets to be launched in the first half of 2010. Keep yourself updated at www.bada.com.

SixthSense to go open source!

Pranav Mistry, a 28-year-old research assistant at the Massachusetts Institute of Technology's Media Lab, on November 8 demonstrated the latest version of SixthSense at the TEDIndia conference, in Mysore. Mistry is the genius behind SixthSense, a wearable gestural interface that augments the physical world around us with digital information and lets us use natural hand gestures to interact with that information. Mistry now plans to open source his project in less than a month.



The SixthSense prototype comprises a pocket projector, mirror and camera worn in a pendant-like mobile device. Both the projector and the camera are connected to a mobile computing device in the user's pocket. The system projects information onto the surfaces and physical objects around us, making any surface into a

digital interface; the camera recognises and tracks both the user's hand gestures and physical objects using computer-vision-based techniques.

SixthSense uses simple computer-vision techniques to process the video-stream data captured by the camera and follows the locations of coloured markers on the user's fingertips (which are used for visual tracking). In addition, the software interprets the data into gestures to use for interacting with the projected application interfaces. SixthSense also supports multi-touch and multi-user interaction. The current prototype system costs approximately \$350 (around Rs 15,000) to build.

Nokia finally starts shipping N900

The highly anticipated Maemo Linux-based N900 smartphone has officially started shipping. The Nokia N900, which has generated a lot of interest since its public launch in August, runs on an ARM Cortex-A8-based OMAP3430 system-on-chip with 1 GB of memory and 32 GB Flash memory. It offers 3G HSPA data and voice, and includes a 3.5-inch (8.9 cm), 800 x 400 touch screen and a 5 mega-pixel camera. It has a QWERTY keyboard, plus Wi-Fi, GPS, Bluetooth and accelerometers.

The feature-packed N900 is going to find its place among the latest high-end smartphones. Its competitors are iPhone, Samsung's LiMo-based Vodafone 360 H1 and Motorola's much-touted Android-based Droid.



First Android phone from Sony Ericsson

Sony Ericsson has launched its first ever Android-based handset. The XPERIA X10 introduces a new UX platform that will evolve across the product portfolio and expand over time, continuously introducing new features and capabilities. The UX platform builds on top of Android and combines entertainment features with signature applications, integration of social media services and a graphical user interface. The X10, which will be available in Q1 2010, features a 4-inch-wide (10.1 cm) touch screen, uses Qualcomm's Snapdragon processor, and comes with an 8.1 mega-pixel camera.

Introducing signature applications like Mediascape and Timescape, the XPERIA X10 lets users organise everything on their phone in an intuitive way. Sony Ericsson Timescape manages all your communications with one person, in one place. Browse through your conversations on Facebook, Twitter, photos, e-mails and texts, all in one go. Sony Ericsson Mediascape claims to be the smart way to get all the music, photos and videos you want from your favourite friends and artists.

We are the Authorized Training Partners For...



Code less.
Create more.
Deploy everywhere.



VMware vSphere 4 : Install, Configure, Manage [V4]	4 Days
VMware vSphere 4 : What's New [V4]	2 Days
VI3 Fast Track	5 Days
VMware View 4	3 Days
VI3 Install & Configure	4 Days
VI3 Deploy Secure & Analyze	4 Days
VI3 VMware Overview	1 Day
VI3 Virtual Desktop Manager	4 Days
VMware Plan & Design	2 Days
VMware Capacity Planner	2 Days
VMware Virtualization Assessment	2 Days

Qt Programme	Ver 4.5	4 Days
Qt Programme Advance	Ver 4.5	5 Days
Qt for S60		2 Days
Qt Embedded		5 Days

Symbian C++ Basics	3 Days
S60 Platform C++ Development	4 Days
S60 Platform 3rd Edition	2 Days

C/C++	Linux Administration
Java	Linux/Windows Device Drivers
Storage Management	JBoss
Linux Basics	Software Testing

December 2009 Training schedules

☞ VMware vSphere 4 : Install, Configure, Manage [V4]

- 01st to 04th Dec'09 Hyderabad
- 08th to 11th Dec'09 Bangalore
- 15th to 18th Dec'09 Mumbai
- 21st to 24th Dec'09 Bangalore

☞ VMware vSphere 4 : What's New

- 29th to 30th Dec'09 Chennai
- 28th to 29th Dec'09 Gurgaon
- 30th to 31st Dec'09 Bangalore

☞ VI3 Deploy Secure & Analyze

- 15th to 18th Dec'09 Bangalore

☞ Capacity Planner

- 28th to 29th Dec'09 Bangalore

☞ Qt Programme : Qt Programming Ver 4.5

- 15th to 18th Dec'09 Bangalore

☞ Symbian C++ : Basics

- 03rd to 05th Dec'09 Bangalore

☞ S60 Platform C++ Development

- 09th to 12th Dec'09 Bangalore

January 2010 Training schedules

☞ VMware vSphere 4 : Install, Configure, Manage [V4]

- 12th to 15th Jan'10 Bangalore
- 19th to 22nd Jan'10 Chennai
- 27th to 30th Jan'10 Bangalore

☞ VMware vSphere 4 : What's New

- 28th to 29th Jan'10 Bangalore

☞ Qt Programme : Qt Programming Ver 4.5

- 19th to 22nd Jan'10 Bangalore

☞ Storage

- 27th to 29th Jan'10 Bangalore

For More Details Contact us:-

madhusudhan@gte-india.com : +91- 984 522 7221
 manjunath.kiran@gte-india.com : +91- 988 022 3289
 basavanna@gte-india.com : +91- 963 236 6332
 bharath@gte-india.com : +91- 963 236 6332
 training@gte-india.com : +91- 80-2669 5890-94[5 lines]

:- Upcoming Training modules by 2010 :-

- Qt Maemo
- Qt Programme Ver 4.6
- Qt Jambi
- S60 Platform 5th Edition
- Flash
- Web Run Time
- VMware Service Delivery -Plan & Design

TRAINING LOCATIONS :- • Bangalore • Chennai • Gurgaon • Hyderabad • Mumbai • Mysore • New Delhi • Pune

For Further Details, Please Contact



G.T. Enterprises

"G T House" #48, 1st "B" Cross, 7th Block, Bhavani Layout, Banashankari 3rd Stage, Bangalore - 560 085.

Tel: +91-80-26695890-94 (5 Lines), Fax: +91-80-26695887 Email : training@gte-india.com, URL : www.gte-india.com



The Ultimate Distro Showdown

It's (K)Ubuntu 9.10 vs openSUSE 11.2 vs Mandriva 2010.

We laid our hands on all the three biggies—Ubuntu 9.10, Mandriva 2010 and openSUSE 11.2—and pitted them against each other. What followed was the battle of the century, as each distro pulled off one unique trick after another to stay on top of the game.

Ladies and gentlemen, behold, as we bring out the next generation of Linux-based operating systems, and try to find out which one has that little extra to make the cut as the best of the best!

The hardware

- *Processor:* AMD Phenom X3 8650 @ 2.3GHz
- *Chipset:* Nvidia MCP67 (nForce 630a)
- *Graphics Processor:* Nvidia GeForce 9400GT with 1GB of GDDR2 RAM
- *RAM:* 2GB of DDR2 800MHz Transcend JetRAM
- *Hard Disk:* Seagate Barracuda 320GB/Western Digital Caviar 160GB (hosts the OS)
- *Screen:* 15" TFT with a resolution of 1024x768

That's a pretty respectable multimedia PC, except for the rather small screen size. However, it shouldn't matter, because if the UI looks cluttered on my monitor, then the UI design is bad. Anyway, with the hardware specs behind us, it's time to move on to the first distro of this review—Ubuntu 9.10.



Ubuntu 9.10—The Koala screams for attention

The Koala may have been the most publicised and talked about release of this month, something like the Linux 2.0 (a la the Web 2.0) for home users, aiming for a level of polish in its user interface hitherto unseen in the world of GNU/Linux. Apart from features like KMS and GRUB2, Mark Shuttleworth had considered ditching the yellow-brown colour palette for something new. Changes like ditching Pidgin for Empathy, which were rather radical at the time of the announcement, dismayed many Ubuntu fans.

Yet, Ubuntu has made it big. The user interface, though a bit loud in its shade of lemon yellow, is laudable

(see Figure 1). It's the same old GNOME though, and at version 2.28, a bit cleaner than the previous releases. In fact, I don't know if it's just me or it's the Ubuntu customisations, but I did stumble a bit after the first boot into GNOME 2.28.

Continuing about the downside of Ubuntu 9.10, KMS and GRUB2 also came as big disappointments to me. Since I have an Nvidia GeForce card, KMS didn't work at all (and sample this, an interview on Phronix says that Nvidia has no immediate plans to introduce KMS support into their drivers) and that, along with GRUB2, made for a very unpleasant boot indeed. It seems that the developers put so much effort into developing the KMS experience that they neglected the non-KMS sector. And considering that KMS is fully supported only on Intel chipsets, and a lion's share of users use Nvidia and ATI cards, that's a very bad decision indeed. Even Fedora was better than this.

Now for the good things... Empathy. Pidgin, heralded as the irreplaceable chat client in the world of both Linux and Windows, has finally met its match, and it seems our little bird flew away to the Empathetic competition. Empathy has come a long way since its introduction and is now a pleasure to use—and it's certainly better than Pidgin, considering its support for rich chat features like voice and video.

The next unique offering is the Ubuntu Software Centre (Figure 3). Rather than use Synaptic by default for package management, it's hidden away for the advanced users. For the layman, the Ubuntu Software Centre acts rather like an app store (except that there's only free software here). Once you open it up, there's a screen with lots of sections. Let's say you want to install Amarok, which you know is an audio player. So you click on *Sound And Video*. From the listing, select Amarok. An arrow will appear in the selection, on the right side. Click on it. A mini-review of the software will appear, briefly describing it, along with a screenshot. Below it, there's a button. Click on it to install. Simple, isn't it?

Ubuntu also holds the distinction of being the first distribution to pay attention to those finer details that differentiate the polish of, say, Mac OS X and Windows to the crudeness of Linux. Packing in more smooth transition effects, extra wallpapers, visual styles and fonts than ever, users are now spoilt for choice (but weren't we before?) over how to make their system look. Coming to looks, Ubuntu has now adopted the Dust colour scheme in the Human theme, with Dark Chocolate Brown window borders and selections, and off-white panels and bars with a slight chocolate hue. Nice colour palette, but I'd rather stick with the Dust theme. There's a new icon scheme as well, called Humanity, but I don't know if it's just with me but many icons don't appear in the Places menu and none in the Systems menu. Apart from this, Humanity is another revolutionary icon theme with mind-blowing icons. Overall, it's a pleasure to actually work in this distro.




Figure 1: Ubuntu desktop after installation

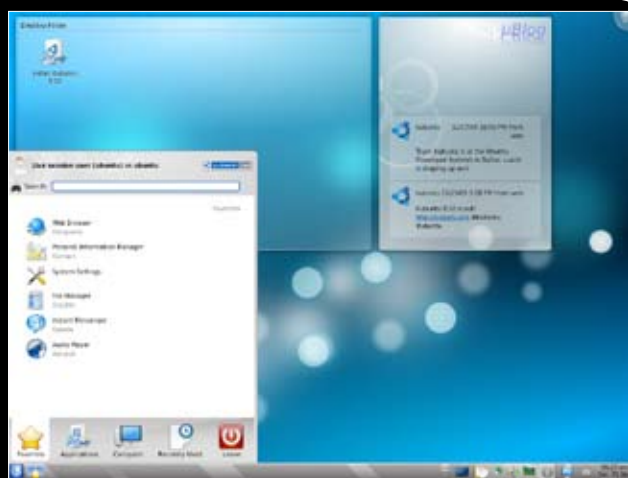


Figure 2: Live Kubuntu desktop with the Kickoff menu and µBlog widget

Now it's time for Kubuntu. At first glance, it's the same boot from Jaunty, and the stock KDE 4.3.2 without any customisations (see Figure 2). But look again and it's a whole new story. It's KMS here too, but this time KMS is a lot more subtle and at least the boot on non-KMS machines looks as good, if not better, than in Jaunty. With KMS on VirtualBox, those transition effects and the flicker-free boot experience did wonders. Boot into KDE4, and the KDM theme just blows you away. Enter your login details and you'll be looking at the prettiest stock KDE4 you have ever seen. And mind you, KDE looks best in its stock theme. Change the looks even a tiny bit and you'll cause a major wardrobe malfunction (or fashion disaster, however you want to put it)!

Kubuntu's installer is another brilliant feature. A full screen image background modelled on the Air theme, with an installer wizard in the centre, guides you through the system installation.

I had a few complaints about Ubuntu, but have none for Kubuntu! None, except for font-rendering issues with

OpenOffice.org, but I'm pretty sure that's because of the fact that OOo is running under Qt. OOo never did like Qt!

OpenSUSE 11.2—The Wow starts now!

Remember January 30, 2007? With a troupe dancing in front of the Taj Mahal, Microsoft Corporation released Windows Vista worldwide on that day. The catch line for the OS? "The Wow Starts Now!". Back then, the company was justified in saying so. But not any more. Enter openSUSE 11.2 and you are sure to get blown away! Let's see what has been done...

KDE 4.3, to start with. Yes, the team has actually gone ahead and customised KDE4, and what a job it has done! The default Air theme has been used, with the thinner panel (this was what actually made me switch to KDE4 full time), and the thick obnoxious green has been replaced by a lighter moss shade (see Figure 4). In fact, the wallpaper is just the default Air wallpaper painted moss colour, with Air-like discs drawn over it. Wow!

Frankly, GNOME 2.28 looks new (see Figure 5). The problem with Ubuntu was that you felt like you were using the same old GNOME. With openSUSE, GNOME doesn't exactly feel like GNOME. It feels better. Everything, from the icon scheme, the dark taskbar, the wallpaper, and the single-panel design, to the custom launcher and the colour scheme, just gels together and creates one hell of a user experience. They call it Sonar. I say 'Wow'. This is GNOME at its best.

The other stuff is SUSE-typical. No sub-pixel rendering support in freetype, Cairo and XFT makes fonts look ultra-ugly—the fix is mentioned later. No multimedia support by default—the fix is mentioned later.

On the software front, they seem to have done some homework. SocioNet now has a privileged place in the operating system, with Choqok and Gwibber, two new micro-blogging clients, of which the latter now supports Facebook. Kopete also now supports Facebook Chat (Oorah!) and new plasmoids are on the cards for Twitter,




Figure 3: Ubuntu Software Centre—a one-stop app store



Figure 4: Default openSUSE KDE with their customised Air theme



Figure 5: openSUSE's GNOME with the brand-new Sonar theme

Identi.ca and openDesktop.

On the technical front, you can now use BtrFS as your root partition, since it now has a stable disk layout! Other than that, the new partitioner in the installer does improve usability a lot. Coming to YaST!, it seems to have stabilised a lot, and for the first time did not crash on my computer while installing openSUSE! Zypper can now download all packages before installing them, and GNOME has an updated Software Update application. And the new artwork stack is, again, Wow!

For netbookers, "...special attention has been paid to the netbook market, and new drivers will enable openSUSE to run better on a wider range of netbooks," says the website. Cool, uh?

Just an aside: When you first install OS11.2, your desktop may be unresponsive and you may experience massive disk thrashing. (This happened on my computer.) This is caused due to an anomaly in Beagle Desktop Search. To fix it, just switch to the console via Ctrl-Alt-F2, log in as the root user and run 'reboot'. Seriously, *that* fixes it.

Mandriva 2010.0—OMG! What's this?

It's a visual disaster—nothing has changed since the last release. Nothing at all! la Ora, the artwork stack of Mandriva (for... what, the last five decades?) has not changed and still shows the thick unbearable purple and blue colours. It's a look that actually reminds you of the golden age of computers—the 90s.

Mandriva does not fall behind when it wants to use the latest and the greatest software (KDE 4.3) but can't seem to resist making it behave like KDE1 either (no widgets, older style menu and panel, folder view and that obnoxious visual style—refer to Figure 6). So, ladies and gentlemen, what is the use of using KDE 4.3 if you are going to disable desktop widgets and fancy graphics anyway?

GNOME? Still the same desktop from the last four releases (Figure 7). No changes. Yet, it's the latest version, 2.28. Wake up, Mandriva devs! I thought you guys were interested in modernisation when you helped K3b port to Qt4! Quoting the *2010 Reviewer's Guide* at wiki.mandriva.com.




Figure 6: Mandriva's KDE4 tries to imitate KDE3




Figure 7: Mandriva's GNOME uses the same old la Ora theme

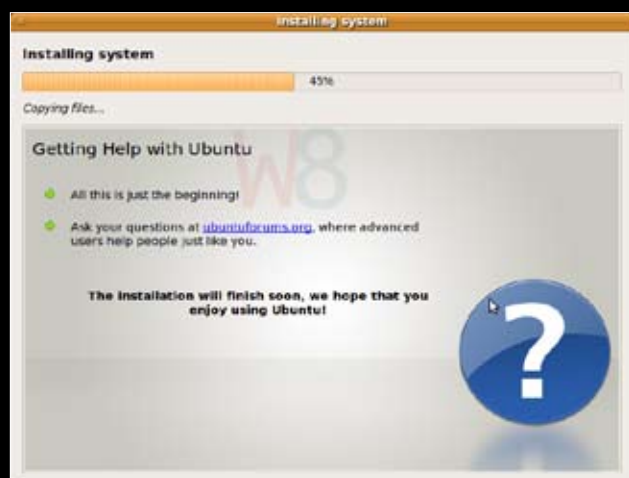


Figure 8: Improved Ubuntu installer



Figure 10: Same-old Mandriva installer

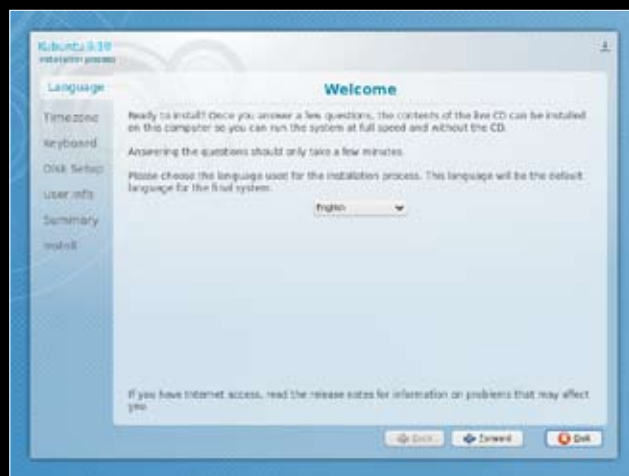


Figure 9: Brand-new Kubuntu installer



Figure 11: Redesigned openSUSE installer

com/en/2010_Reviewers_Guide: “The big, obvious, can’t-get-away-from-it new feature for 2010 is obviously KDE4 as the default desktop...The point is, we think our implementation is probably the best among this cycle’s major distribution crop. So go ahead and compare it with Kubuntu, openSUSE and so on.” They can actually say that after making KDE4 look like goo?

Mandriva does still score though. It’s the only distro of the three to actually include support for patent-encumbered multimedia codecs, out-of-the-box. Mandriva’s Control Centre is another innovation that heavily tips the scale in its favour, and the large amount of software included on the DVD means you don’t need an Internet connection for anything except updates. Ditto for one gazillion games, and one quintillion software packages—everything from HPC clusters to mundane office suites and everything in between. Cool, huh? And surprise of surprises—Mandriva, by default, includes Google Gadgets For Linux (GGL) for GNOME users. What? Disable KDE’s

widgets by default and give GNOME users GGL? Is this some sort of revenge?

Although GStreamer and Xine plug-ins are available on the DVD, playing a MP4 downloaded from YouTube in Totem on KDE4 (that’s the default) got stuck because an AAC codec could not be found. Rather than download free codecs or ask me for the DVD, Codeina insisted I purchase a set of licensed codecs from Fluendo, for 11 Euros. Strange!

Okay—review’s over. Time for a shoot-out.

The distro shootout

It’s time we pitted these distros against one another. They will be checked on the basis of their download package, install experience, included software (as installed on the discs as extras) and software versions, looks, multimedia support, package managers, configuration centres and customisations, besides unique features. There are no marks or points in each category; it’s just *The Good*, *The Bad* and *The Ugly*.

Criteria 1: The packages

Mandriva as usual comes on a gigantic 4.4GB DVD image, filled up to the brim with software. The inclusion of every possible software package imaginable means you won't have to reach for your Internet connection in any but the most exotic of conditions. openSUSE is another DVD, albeit a bit smaller at 4.2 GB. It does not include as much software, but has something called 'Installation From Images'—actual disk images in which patterns of software are pre-installed from the RPMs and don't have to be installed manually—and this speeds up installs considerably. Ubuntu and Kubuntu are LiveCDs, each weighing in at 699 MB, and although the smallest download of the lot, it's unusable without an Internet connection because additional software needs to be downloaded from the Net before you can work on the OS.

Judgement: Unjudgeable. All of them are unique. openSUSE has a good blend of software on DVD, Ubuntu has an uber-small download size, and Mandriva includes everything but the kitchen sink on the DVD. It's your personal taste.

Criteria 2: Installation experience

In Ubuntu, the Ubiquity installer seems to improve drastically, release by release. Apart from being more stable, it now has even more polish. A slideshow during the set-up phase (Figure 8) actually enlightens us about the practical features of the operating system rather than brag about how this particular OS is the best of the lot (as openSUSE does) or advertise other products by the company (as Mandriva does). Out at the Kubuntu stable, the installer is one H-U-G-E surprise, with such a completely redesigned minimalist yet pretty UI (Figure 9) that you'll actually want to stare at it rather than do anything on it. I personally think this actually compares with the likes of the Windows 7 installer.

Over at Mandriva's stable, it's again the kitchen sink installer, and this time they haven't even changed the UI colours or theme (Figure 10). As a result, like the rest of the system, this looks ugly. There is a lot to configure in the installer, including installing everything minus the server tools, LXDE and XFCE, and about 30 MB worth of updates; so it took more than an hour to complete. Sure, if you choose to install only KDE4 and no updates, you can finish it in 15 minutes. But anything more and you are done for.

openSUSE is a very pleasant experience. The installer is darn pretty (Figure 11) and does as much as the Mandriva installer can do, but in a much better and more understandable way. The layout of the installer—presenting the installer pages as rich pages rather than dialog boxes—radiates charm. The 'Installation From Images' feature speeds up installation. And if the installer finds an Internet connection during

installation, it automatically installs Adobe Flash and Microsoft TTF fonts. It can't install the codecs, though, because those stay in the third-party Packman and VLC Repos. It's one of the best-designed professional installers I have ever seen.

The Good	openSUSE 11.2, Kubuntu 9.10
The Bad	Ubuntu 9.10
The Ugly	Mandriva 2010. And it really is ugly

Criteria 3: Software

Mandriva fans rejoice! I have nothing more to say. Want KDE4? GNOME? LXDE? XFCE? Games? Servers? HPC Clusters? Databases? Three office suites? All it took me was 10 minutes to download the Nvidia drivers, and I was playing *Battle For Wesnoth* on Mandriva before I knew what I was doing.

All three have Kernel 2.6.31, so I can't pull anything here.

The Good	Mandriva 2010
The Bad	openSUSE 11.2
The Ugly	Ubuntu 9.10

Criteria 4: Looks, GNOME and KDE4 implementations

Does it even look like openSUSE has any competition here? Had they included the sub-pixel rendering patches for freetype, XFT and Cairo, openSUSE with its lime green KDE4 (Figure 4) and dark suede green GNOME (Figure 5 and Figure 12) would have scored a perfect 10 in the looks department. Even without them, nobody comes even close. This has the best KDE4 implementation around.

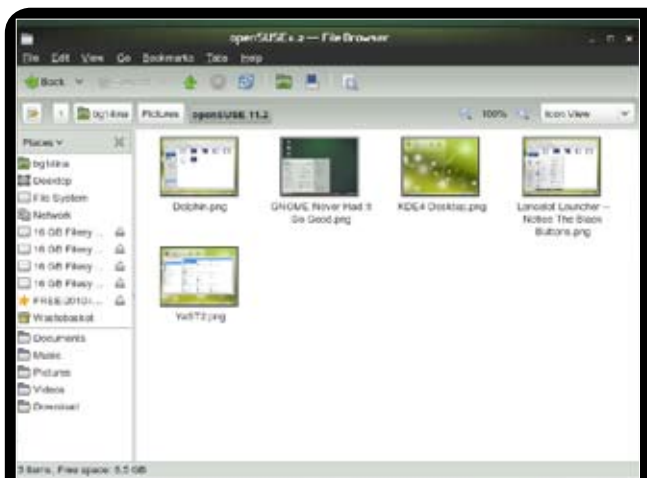


Figure 12: Up, close and personal with the Sonar theme using Nautilus

Ubuntu hasn't let go of the old GNOME that paints the whole town yellow (Figure 1). Brown and yellow have now been used for too long, and that bright yellow wallpaper gets on your nerves. However, I've got to admit, they have the best GNOME implementation around. That sounds contradictory? Well, looks are not everything; what counts is how integrated the DE is with the rest of the system and vice versa. If you want GNOME, nothing else will do. Kubuntu is a refreshing surprise though, and it's not just KDE4 that looks good (Figure 2); Kubuntu has perfected the desktop layout, text and panel sizes and text anti-aliasing—the results are awesome.

Mandriva makes me angry. And I'm sure many Mandriva users will be, after checking out its hideous looks. I'll forgive you if you spread the news that Mandriva uses KDE1 as its default UI (Figure 6). I'll forgive you if you think this release was made in 1989. On the upside, the Mandriva developers have worked on KDE's Semantic Desktop (the Nepomuk-Strigi combo) feature and, unlike the other two, it's enabled and integrated out-of-the-box. Although I'm still not used to the idea, you can have a better understanding of its capabilities at <http://doc4.mandriva.org/bin/view/labs/Nepomuk-mdv2010-RC>.

The Good	openSUSE 11.2
The Bad	(K)Ubuntu
The Ugly	Mandriva 2010

Criteria 5: Multimedia support and other proprietary stuff

Let's start with Ubuntu 9.10. The only codecs that it comes with are Ogg Vorbis and Theora. However, installing additional ones is a breeze. Once you play any

unsupported file (such as MP3) in Totem or Rhythmbox, a dialogue box pops up and installs all the GStreamer codecs. Flash, however, needs to be installed manually. Over at Kubuntu's stable, it's much better, as Kubuntu automatically installs all the codecs, including Flash, the moment it detects a working Internet connection. So one doesn't really feel that any codecs are missing! Way to go!

Coming to openSUSE 11.2, even this guy comes with only the Ogg codecs. However, no automated installation procedures exist. What you need to do is get into YaST!, open the Software Repositories module and hit Add. In the next screen, select Community Repositories and enable VideoLan and Packman repos. Finish up the wizard and the Software Repositories module, and close YaST!. Open a terminal, and run: `sudo zypper dup`. You are good to go.

And at Mandriva's stable, everything is included, except very exotic stuff like Dirac and AAC. You can just head to Mandriva's Control Centre and then to the Package Manager; then search for `gst`, and enable everything. You are covered—mostly! Add the PLF repos (from easyurpmi.zarb.org), open a terminal and execute `urpmi --auto-select`. This will replace the somewhat restricted codecs with the fully-liberated ones—thanks to the Penguin Liberation Front :-)

Coming to proprietary graphics drivers, over at Mandriva, this is one no-brainer to install. Head to MCC, and then to X configuration. Don't change anything; just keep hitting 'Next'. Just before the summary screen, MCC will inform you that proprietary drivers that perform better are available, and if you would like to use them, hit 'Yes', and they are installed. Log out and back in, and it's done.

In (K)Ubuntu, the Jockey applet starts up on the first boot, showing a list of hardware that need proprietary drivers (it included only Nvidia G96 (9400GT) in my case)—refer to Figure 13. Select a driver version, hit 'Activate', and reboot. (You don't really need to reboot;



Figure 13: Ubuntu suggests I should be better off with NVIDIA's non-free drivers



Figure 14: The Mandriva control center continues to make life easier



Figure 15: A redesigned Qt4 version of YaST2 UI

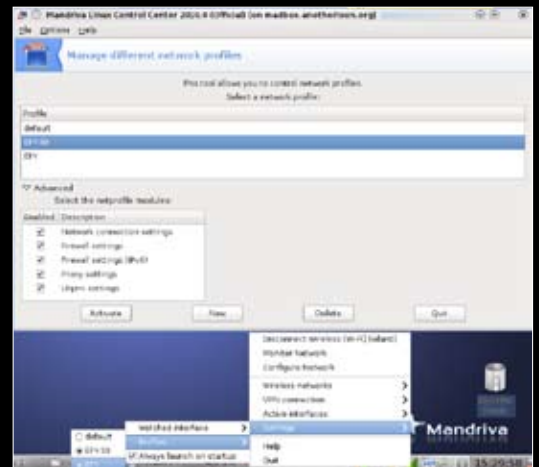


Figure 16: Network profiles are a life saver when you're on the move

open a terminal, run `sudo modprobe nvidia`, log out and back in to X, and you are done.) Simple.

And in openSUSE, you need a Web browser. Head to <http://en.opensuse.org/NVIDIA> or <http://en.opensuse.org/ATI> and then use the one-click feature to install the driver. Simple again!

The Good	Mandriva 2010
The Bad	(K)Ubuntu 9.10
The Ugly	openSUSE 11.2

Criteria 6: Package manager

Ubuntu uses DEB packages, while Mandriva and openSUSE use RPM. For those not familiar with these terms, DEB is an unparalleled package management system that just has no match. Need I say more?

Yes. Between the two, I found *zypper* easier to use compared to *urpmi*, as the former's options to install and remove packages are similar to that of *apt-get*. Besides, Zypper includes an option to download all RPMs before installing them. However, it's not enabled by default, and you must edit the `/etc/zypp/zypp.conf` file and add the following line under the relevant section:

```
commit.downloadMode=DownloadInAdvance
```

That should do the trick!

The (very) Good	(K)Ubuntu 9.10
The Bad	openSUSE 11.2
The Ugly	Mandriva 2010

Criteria 7: Config centres

Mandriva has something called the Mandriva Control Centre (Figure 14). Windows users, who have Control

Panel, will finally say Linux has a match in the Control Centre department. From MCC, one can administer every aspect of the system, including, but not limited to, adding/removing software, administering hardware, networks, printing, the boot process, user accounts, the look and feel, and a lot more. And it's neatly categorised with a very slick UI.

OpenSUSE, on the other hand, has YaST!, or Yet Another Setup Tool! (it has an exclamation mark at the end). As you can see in Figure 15, this one is also cleverly designed and a bit more powerful than MCC. However, this is at the cost of user friendliness.

Oh, and Ubuntu and its brothers have no configuration centres at all.

The Good	Mandriva 2010
The Bad	openSUSE 11.2
The Ugly	(K)Ubuntu 9.10

Criteria 8: Customisation, unique features, the 'X' factor, and our beloved bugs

Let's start with dialogue boxes. In openSUSE, applications like Firefox, OpenOffice.Org and others like Thunderbird use KDE4 and GNOME dialogue boxes for Save/Open dialogues based on the DE they are running under. This happens in Ubuntu as well, but not in Mandriva.

openSUSE, as usual, seems to be housing some bizarre bugs. While running under KDE4, OpenOffice.org apps seemed to have no Status Bar; however, randomly clicking in that area did make indicators like INSRT/OVR appear. And back when I first installed and started openSUSE, it froze up for no reason at all. Switching to a text console and running `top` seemed to suggest that Beagle Desktop Search was using up a lot of processing resources. The HDD data activity indicator

Musings

1. Mandriva's drawbacks: ugly looks, occasional unresponsiveness, and the 5353 thing, plus trying to make KDE4 look like KDE3 on the outer surface.
2. GRUB2 included with Ubuntu and Kubuntu makes them un-chainloadable from the Windows boot manager, and can be booted from GRUB legacy -- but this needs a flabbergasting GRUB.CONF.
3. openSUSE 11.2 includes an old KDE4: 4.3.1, while the others include 4.3.2. No updates are available. Even with Amarok, its version 2.1.1, not 2.2—you'll need to pull the latest from the KDE Backports repo, manually.
4. Anti-aliased fonts need to be enabled manually via KDE system settings or Configure Desktop in GNOME, in openSUSE 11.2. Or else they look U-G-L-Y. And the fix is per-user only; KDM and YaST! running as root still display Monochrome fonts.
5. A stupid boot experience for non-KMS Ubuntu users—randomly, the entire boot happened in text mode.
6. Some encrypted home features in Ubuntu cannot be mounted in any other operating system—specifically, openSUSE and Mandriva. It did work in Sabayon Five Oh.

on my chassis was lit up like a Christmas tree. Killing the process made no difference at all, and then I ended up rebooting the PC, and everything was fine again. Mind though, this was RC2; I actually installed RC2 and then updated to finally using a DVD. It may have been fixed (though the Beagle package was not updated). Also, with openSUSE I could not manually mount the ECryptFS'd home from my Ubuntu; it gave me quite interesting filenames.

And Mandriva comes up with a bug now! Every 15 seconds or so, a notification comes up, saying "Linux.local is connecting on the 5353 service". It gets royally annoying. Tracking it down, I discovered that UDP port 5353 is used for multicast DNS (mDNS). I asked the Mandriva fan at LFY (who knows its internals more than he knows his own ;-)) what mDNS was all about, and he pointed me towards Avahi. Disabling it fixed the problem.

Get to networking and Mandriva comes up with one slick feature. Mandriva includes support for something called Network Profiles—refer to Figure 16. Let's suppose you use a wired network at home, Wi-Fi with a corporate proxy at office, and a public profile with full firewall, DHCP and no network daemons in public places. With Network Profiles (accessible from MCC -> Network), you can create profiles of network settings, say one with static IP and no proxies for your home ADSL router; one with DHCP over Wi-Fi, and a corporate proxy for the office; one with

maximum security for public places like airports, and then switch between them at will. Pretty nifty, because at home it saves the overhead of a DHCP request. At the office it makes sure all the proxies and DHCP stuff work homogeneously all over the system, and in public places it maintains some sanitised security. Cool eh?

While we are at networking, openSUSE automatically detects if you are using a laptop or a desktop, and then uses the *ifconfig* on desktops (because they are mostly static) and NetworkManager on laptops. Ubuntu uses NetworkManager as the one-size-fits-all tool. While it works, there are dynamic configuration overheads.

The Good	Mandriva 2010
The Bad	openSUSE 11.2
The Ugly	(K)Ubuntu 9.10

Overall

openSUSE 11.2 left the best impression. Its absolutely killer looks, slick integration, superb tools and really good repository and package count makes it THE choice for all but novices or very finicky Linux users out there, Netbookers apart.

Ubuntu and brother Kubuntu are cool and minimalist. Perfect for people who want DEBs at all costs and want the system customised exactly their way. The light and minimalist nature of the system means it's a perfect choice for netbookers and low-end laptops. However, with only a frighteningly low 5 per cent of all Ubuntu users reporting that they had a flawless install (95 per cent said they had serious problems), it does put Ubuntu in a spot. Coupled with the emphasis on KMS and the resulting neglect of non-KMS users, this makes it to the second spot.

And staggeringly far behind, with its scary looks, its occasional unresponsiveness, a bloated software base and a very inconvenient installer, Mandriva makes it to the third spot. Only for switchovers from Windows (they'll love Mandy) and novice Linux users. Real men use openSUSE 11.2!

The Good	openSUSE 11.2
The Bad	(K)Ubuntu 9.10
The Ugly	Mandriva 2010



By: Boudhayan Gupta

The author is a 15-year-old student studying in Class 9. He is a logician (as opposed to a magician), a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++, Python and Assembly (NASM Syntax).



RHCE | RHCSS



CCNA



SOLARIS



DBA

▶ Redhat Training & Certification Center

▶ Pioneers in LINUX Training

▶ Certified Faculties

▶ Corporate Batches -DELL, NRSA, GENPACT ETC.,

▶ Weekend Batches

▶ Instructor Led Training

▶ 1:1 Systems / Routers

▶ Placement Assistance



Amrita Technologies
(Software Division-Hyderabad)

844/1, Mahindra Hills, (E) Marredpally, Secunderabad - 500 026.

Tel : 9393733174 | 27733174 | 27737969

Email: aict.hybd@amrita.ac.in | www.amritahyd.org

www.eleb2b.com

**Electronic
Components
Products
Machinery
Will now be sold here**

Presenting Electronics Bazaar,
India's First B2B Sourcing Guide
For Electronics Components,
Products and Machinery.

Unit Price: **Rs 50**

Subscription Rate: **Rs 350 (12 Issues)**

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area,
Phase-I, New Delhi 110 020; Ph: 011-26810601-03; Fax: 26817563;
Email: info@efyindia.com, Website: eb.efyindia.com



ELECTROMECH

**RHCE / RHCSS Training
Training by experts**

Linux solutions

(Mail server, Samba server, Proxy server etc.)

Open Source Development

**RHCE exam 27th DEC
Ahmedabad – Gujarat**

"Come home of Open Source"



Redhat Channel Partner and Training Partner

Ahmedabad: 404, Maulik Arcad,
Nr. Mansi cross Road, Satellite Rd
Baroda : 25, The Emperor, Fatehgunj

Ph.: 079 40027898/9327218003

Http://www.electromech.info

E-mail: electromech@electromech.info

FOSS in Action

Preparing for that Crucial Meeting



The FOSS ecosystem gets richer by the day. A tool-based approach of looking at FOSS applications in isolation may not be as convincing as watching FOSS in action in an every-day office scenario.

*T*he general tendency is to evaluate FOSS applications based on particular tools. This leads us to look at FOSS tools in isolation and describe their functionality. The user scenarios are generally taken for granted or are implicitly assumed. This approach works fine most of the time. But when overdone, there is the danger of missing the forest for the trees. A use-case-based approach, on the other hand, allows us to set up a compelling context around the tool's usage and allows us to create a more demonstrative presentation for decision makers. Let us follow Aakash as he convinces his bosses about the superiority of a fairly unknown

FOSS tool, by actually demonstrating its capabilities when it was most required.

The background

Aakash's employer, OS² Systems, was a medium enterprise in the services business. It was important to make every penny count at OS² to stay profitable. Although he had other offers on hand, Aakash had been impressed by his employer's recruitment team on campus. "You will have all the freedom to experiment and try out new things to improve our productivity and profitability," the CTO on the recruitment team had told him. The FOSS aficionado that he was, Aakash saw this as an opportunity to put FOSS to work in a real corporate

Disclaimer: Any resemblance to real characters is purely coincidental and unintentional. The choice of desktop environments and applications mentioned is not meant to be prejudicial in any way. The situations are contrived for the sole purpose of effective communication. Finally, YMMV or your mileage may vary.

setting. And he had no regrets. A few days after joining, Aakash had even convinced the management to start an office productivity project centred around FOSS, under his leadership.

Problems as opportunities

Platitudes apart, that is exactly how it happened. Clients were expected the following week for an 'ideation' session. It was to be a freewheeling meeting with white boards and flip charts instead of prepared OHP slides. The minutes of the meeting had to be as accurate as possible and available almost immediately after the meeting, so that they could be agreed upon and signed off. A lot depended on the professional conduct of the meeting to make a favourable impression on the clients.

A solution of sorts did exist. It was the electronic white board, which had two writing surfaces and could copy the contents of the board on to special paper. The results were far from satisfactory and the copies were always in black and white. Distributing the minutes was a chore. But even that facility went out of order with no spares available. Piyush, the head of marketing, was convinced that the only way to make accurate minutes was to take pictures of the white board and flip charts with his digital camera, and attach them to his minutes. "Not a very elegant approach; but it would have to do," thought Piyush. He did mention his problem to Aakash as the two met briefly before Piyush drove home.

Aakash had almost reversed out of the parking lot when it struck him that there was a way out. He re-parked his car and ran back to his cubicle collecting the keys to the conference room from the reception. There was work to be done...

The goal-driven use case here was straight forward. The company needed a 'cost-effective', 'elegant', and 'reliable' means to capture meeting minutes with a minimal turnaround time.

The Saturday movie and after

It was a relatively lazy Saturday after a long time. Piyush and Aakash planned to grab a quick brunch before retiring to the noon show of 2012 and calling it a day. At least that was the plan. Things did go as planned but Piyush seemed preoccupied throughout. After the movie, instead of driving homeward, Aakash steered his car towards the office. "Where are we going?" asked Piyush. "Wait till I show you something," said Aakash.

Once at the office, they quickly went into the conference room. Waiting there was a laptop with a digital projector and digital tablet connected to it. "Is this some kind of a joke?" thought Piyush, not quite appreciating the set-up; yet he decided to stay quiet and let Aakash speak first.

What we are about to see is one realisation of this use case. There might be others. We deliberately chose a goal-driven use case over one that's 'system interaction' driven to broaden the solution space.

The realisation

The 4 p.m. sun was lighting up the glass panes towards the west. "Lower the blinds, will you?" Aakash requested Piyush. This was getting a little too melodramatic for Piyush's taste. Aakash had already booted into the GNOME desktop and was scrolling down items in the 'Office' menu. The next moment, an application screen filled the far wall of the conference room as the digital projector warmed up. It was a giant notepaper with a couple of tool bars on top. Piyush wished this thing would end or that Aakash would say something. Instead, Aakash picked up the stylus of the digital tablet and wrote 'Get it?'. A huge version of these words appeared scrawled on the far wall.

Suddenly, Piyush straightened up in his chair. This looked useful.

"This is Xournal," said Aakash. "It is a pen-based note journal that can also be used to take notes. Let me



redhat®

TRAINING PARTNER

RHCE / RHCSS Exam Centre

Hundreds of hardcore
Industry Professionals
Vouch for it.



RHCE only @ ADVANTAGE PRO

- Pioneer in LINUX training since 1999.
- Chennai's #1 RHCE training Centre
- 100% results every month.
- Many centum students.
- Also get trained on My SQL - CMDDBA, CMDEV, PHP, Perl, Ruby, Python, Ajax...

Next RHCE/RHCSS Exam
7th & 21st December '09

Don't wait.
Register Now.

Let the success vouch for you!



ADVANTAGE PRO

IT TRAINING DIVISION OF
Vectra Technosoft Pvt. Ltd.
Open Source Academy

Regd. Off: Wing 1 & 2, IV Floor,
Jhaver Plaza, 1A, N.H. Road,
Nungambakkam, Chennai - 34.

Ph : 28263529 / 3530 / 3540
Telefax : 28263527
E-mail : enquiry@vectratech.in
www.vectratech.in

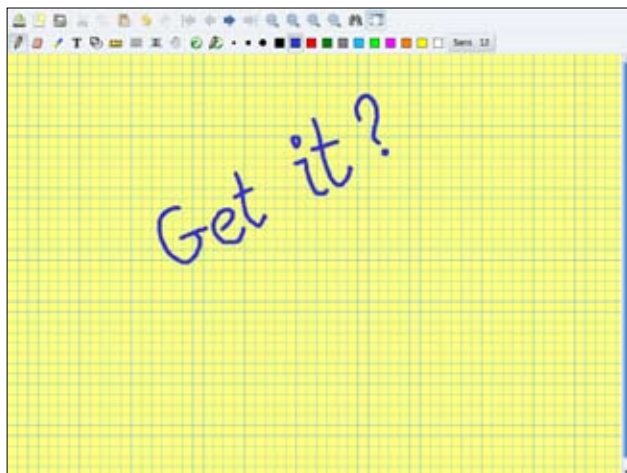


Figure 1: Xournal application

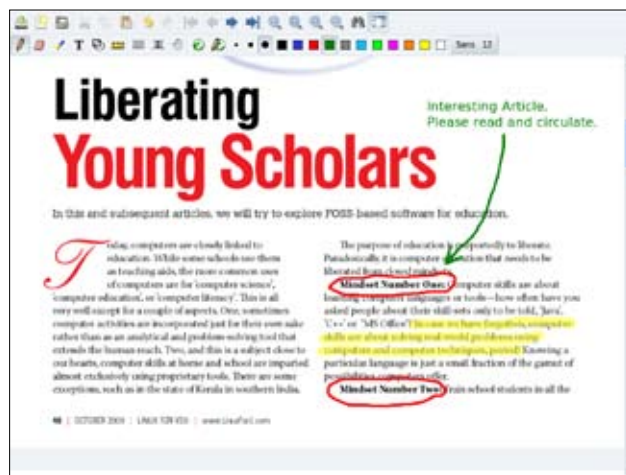


Figure 2: PDF annotation

quickly take you through its features. F11 toggles between the full screen and window mode. You can choose from a variety of backgrounds and note papers. There is even a square-ruled one to help you draw diagrams. You have a choice of pen colours and widths to give clarity to your drawings and notes. You can also insert typographic text at an arbitrary point on the notepaper, allowing you to intersperse hand drawings and typed text. The shape recognition feature converts your squiggled circle into a perfect one. You can even insert a blank space between two sets of handwritten notes and drawings.”

“So, I take screenshots of each screen and...,” interjected an excited Piyush, but trailed off as he heard Aakash saying: “...and you can export this to PDF.” That was it. A projector-friendly application that lets you take free-hand and typewritten notes and export them instantly; he could not have asked for more.

The next hour was spent with an exuberant Piyush rehearsing his role for the meeting by articulating his ideas on a yellow, square ruled paper background with an impatient and hungry Aakash glancing at his watch.

Party time

Monday morning was overcast but Piyush had the sun shining on his face as he strode into the conference room with his clients. The meeting

was very productive with members of the client team taking turns to articulate their ideas and then seeing it appear on the OHP screen. The notes were getting typed on the virtual notepad (remember the typographic text feature) right next to the handwritten block diagrams and flow charts.


The conduct of the meeting was so convincing that an agreement was reached on the minutes, even as they were getting typed. Piyush loaded the exported PDF into his thumb drive as he drafted a formal mail to his clients with the minutes of the meeting as an attachment. As he left the conference room, he overheard one of his clients ask the CEO, “What is that set up you used for this meeting?” The CEO, in turn, looked at Piyush and said, “Remind me later to ask you how you did this...,” as he took the clients to lunch.

The sequel

A couple of days later Piyush was busy scanning some annotated sheets to send out as attachments. They were the marketing team’s comments to a proposal from the branch office.

“Did I tell you that you can use Xournal to annotate a PDF? Just load the file as background via the ‘Journal’ menu, annotate it, and

export it back as a PDF. And you can even annotate files as a meeting activity,” said Aakash as he headed for the door. “Makes it so convenient to use with intermediate work products,” he continued. Before Piyush could react, Aakash was out of the door and had disappeared into the evening sunshine. “Should remember to get Xournal installed for all staff members with review responsibilities,” thought Piyush as he put another page into the scanner. He remembered that OpenOffice.org could export any document to PDF.

Aakash was driving home in the bright evening sun. It was not often that he got to go home when the sun was still up. His project had finally got a higher budget allocation after the CEO had heard of the successful client meeting. He had loaded a couple of new albums on his music player. There was an interesting football match on TV in the evening. Life looked good. **END** 

References

- <http://xournal.sourceforge.net>
- Jarnal [<http://www.dklevine.com/general/software/tc1000/jarnal.htm>] is another annotation program that also supports collaboration.

By: Gurudutt Talgery

The author is an open source enthusiast and practitioner with over 20 years of industry experience working with reputed transnational companies. You can contact him at gtalgeryATrediffmailDOTcom.

Can you meet your email retention challenges?



Emergic MailArray solution makes your enterprise litigation-ready by ensuring on-demand discovery of emails. It optimizes the mail server performance by freeing up primary storage capacity. It enables you to preserve your email in a single, scalable repository and meet the needs of different stakeholders within your enterprise.

Key Benefits

- Compliance (SOX, HIPPA)
- Policy, Security & Retention rules
- Indexing Content & Attachment
- Compression & De-Duplication
- Quick, Easy Search & Retrieval
- Insurance against Litigation
- Encrypted Backup

Partial Client List

- Geojit Financial Services Ltd.
- IDBI Bank
- IDBI Capital Ltd.
- General Insurance Corporation (GIC)
- IndusInd Bank Ltd.
- Apollo Sindhoori Capital Investment Ltd.
- Flat Products Equipment (I) Pvt. Ltd.



sales@netcore.co.in | www.netcore.co.in/ema

MUMBAI	NEW DELHI	CHENNAI	BANGALORE	HYDERABAD	PUNE	AHMEDABAD
9820032324	9350101237	9380388685	9341609111	9392224001	9373727934	9824192328

Netbooks

Facts, Figures, Options & Opinion



It has been a long while since the first laptop made its way on to store shelves. What led to the invention of a laptop was the necessity of a portable computer that could be lugged around and also worked on, in the absence of direct power supply. All went well until mobile phones with Internet connectivity came along, and that changed the whole dynamics of mobile computing.

The mobile Internet, coupled with Internet hot-spots, fuelled the craze of 'anytime, anywhere computing' and suddenly we found our traditional laptops very heavy, cumbersome and in need of more electrical juice. And then, increased bandwidth and more powerful Web apps gave people the possibility of fulfilling all their computing needs online. Thus emerged a new breed of computers, christened 'netbooks'.

What is a netbook?

In Wikipedia's words, "...netbooks are small, light and inexpensive laptop computers suited for general computing and accessing Web-based applications; they are often marketed as 'companion devices,' that is, to augment a user's other computer access."

A typical netbook ranges from 8 inches (20.3 cm) to 11.6 inches (29.5 cm), with smaller keyboards, reduced computing power, fewer ports and no optical drives. Because of their small size and low computing power, the normal operating systems that we are so accustomed to in our daily use, aren't the most ideal. This gave rise to the need of customising the existing operating systems so

that they fitted into a netbook's size and form, without compromising too much on the performance.

What are the operating systems that run on a netbook?

UI designers need to be given utmost flexibility and importance, when it comes to designing an operating system tailor-made for a netbook. Some of the challenges faced while building an OS for a netbook are:

1. The screen size is much smaller compared to a traditional 14-inch (35.6 cm) notebook.
2. The display resolution of a netbook hardly ever crosses 1024 x 768.

The Atom processor has done wonders in increasing the battery life of netbooks. This, coupled with the smaller display, and the way an operating system is coded, play a big role in optimising the battery life.

Windows XP had been shipped with almost every netbook since the very first of the breed ever left the factory. However, a lot of manufacturers are now gungho about the much-optimised Windows 7 and are hurrying to bundle the OS with their machines.


But then, how about the world of open source?

Yes, we are getting to it. :-) It is in the world of open source that most of the operating systems, that are tailor-made for netbooks, have emerged. While Ubuntu had always been in the lead, with the Ubuntu (and Kubuntu) Netbook Remix, Intel is also inching forward with its much celebrated Moblin (recently out of beta and now launched as v2.1). Then there is one of the most interesting distributions we have come across—JoliCloud, developed by Tariq Krim (founder of Netvibes), and based upon Ubuntu Netbook Remix.

What's more, even Google decided not to be left behind and announced its spanking new Chrome OS. We also got to know that Acer has started to offer its Aspire One D250 with dual booting options—Windows XP and Google's Android.

The netbook industry is growing in leaps and bounds. Today, almost every social media entrepreneur sports a netbook. It's the best time for a year-end wrap-up of Linux-based operating systems and also the netbooks that are on offer, today.

However, on a note of regret—while we were promised netbooks from Lenovo, Sony, Acer, ASUS, Dell and Samsung, we finally got trial models only from Acer and ASUS. The rest decided to back out at the last moment since they wanted to test their products with the newly released Windows 7 before they sent them out for review. The fact that we informed them we didn't even plan to boot Windows on their machines, didn't help matters. We apologise for not being able to come up with detailed reviews of other netbooks in the market.



Acer Aspire One D250

Acer made a significant impact on the netbook market last year with its hugely popular Aspire One A110. It recently released a selection of new netbooks in its Aspire One range. The D250 we have here might not be a catwalk queen, but it looks attractive enough and benefits from a well-built chassis that doesn't creak under pressure.

Design

I was more impressed by the D250's girth—or rather, the lack of it. It's only an inch high. That's not as thin as ASUS' svelte Eee PC 1008HA—but is still much nicer than lots of other, chunky netbooks.

The D250 I reviewed makes use of Intel's 1.66GHz Atom N280, and Acer supplies it with 1GB of DDR 2 memory. Watch out, though: it's also available with the 1.6 GHz N270.

Ports

Two USB ports and a multi-format card reader are found on the right, along with the power socket. On the left sit the Ethernet, VGA, audio in/out and a third USB port. The vent is also on the left, and although the D250's fan fired up a fair amount during testing—even when doing little more than browsing the Web—it's not loud enough to irritate.

The battery

The rear is home to nothing but the three-cell 2200mAh battery. A six-cell version is also available, though this battery's a bulky affair that angles down from the notebook to lift the back of the machine up off the deck. It makes the machine less comfortable to carry—stick it in a bag so that it doesn't bother you. You may find that it angles the keyboard better.

At the front, you'll find a small Wi-Fi power switch—Bluetooth can be axed using a button just above the keyboard.

Storage

Flip the D250 on its belly and three removable backplates are revealed, providing access to the hard drive bay (home to a 160 GB 5400 rpm SATA drive), the memory and an empty Mini PCI Express slot. If you want to improve on the memory, you'll need to jettison the installed 1GB module, since the D250 only houses a single SO-Dimm slot.

Keyboard and touchpad

On the 8.9in A110, the keyboard took up almost the entire width of the chassis, barring a 1cm edge on either side of it. Admittedly, on a netbook there's not much room but even a slight enlarging of the keys would improve usability. Acer could also have switched to the flat-style of keys as found on its new Aspire One 751H netbooks. As it stands, the keyboard is a bit too fiddly to type on.

The trackpad measures just 50 x 30mm, and Acer has managed to furnish it with multi-touch capabilities, though these are fairly limited. You can zoom in and out by using a two-fingered pinching motion, while dragging two fingers left and right will perform functions such as navigating back and forth between Web pages and scrolling through an album of photos.

Display and networking bits

The 10.1-inch (25.6 cm) screen has a native resolution of 1024 x 600, and is bright, crisp and exhibits vibrant colours. Best of all, though, it performs amazingly well outdoors and even in direct sunlight I was able to see the desktop clearly. The only thing you'll need to watch out for when using it outside is the glossy screen attracting troublesome reflections.

Located just above the screen is a webcam, but it's only a 0.3 mega-pixel job.

The fact that Gigabit Ethernet has been sacrificed for bog standard 10/100MB/s wired networking won't concern too many prospective buyers, but Acer's decision to go for 802.11b/g wireless instead of 802.11n is surprising to say the least. Bluetooth is built-in.

Verdict

The Aspire One D250 weighs just 1.1 kg, making it one of the lightest 10-inch netbooks you'll find. Taking its lightweight nature into account, the fact that it managed to last two hours and twenty-four minutes through our video playback test, is admirable. You'll get even more back-up by going easier on the D250, and with a fully-charged battery, JoliCloud estimated there was enough juice for about three and a half hours of action.

Speaking of performance, when troubled with a bit of Gaussian Blur action in The Gimp, it took an average of 4.8 seconds to apply the filter to our standard test image.



▲ The hinge connecting the monitor and the keyboard makes the body look sleek



▲ The input ports are intelligently flushed within the body

Acer Aspire One 751H

In many ways, the Acer Aspire One 751h, the company's first 11.6-inch (29.5 cm) netbook, is an improvement over the D250 I had reviewed. With Aspire One, users get an extra inch of screen real estate, a roomier keyboard, and more than five hours of battery life (on Ubuntu Netbook Remix). But although it's larger, the 751h has a weaker processor. The result is a netbook that looks and feels great but offers a slower performance than other machines in its class.

Design

If you've seen the 10-inch Aspire One D250, the 11.6-inch version will look familiar. It also has a glossy black lid with Acer's metal logo in the lower corner; the lid picks up fingerprints fairly easily. There's also a grey palm rest and keyboard deck with a black keyboard and thin bezel to match.

Thus, the overall effect is conservative, but not cheap-looking. The real cornerstone of the Aspire One's design is not the colour scheme anyway, but its thinness. Whereas its competitors, such as the ASUS Eee PC 1005HA, are 1.4 inches (3.5 cm) thick and weigh close to 1.5 kg (the 1005HA weighs 1.4 kg), the Aspire One has a minimal footprint of 11.2x7.8x1.0 inches (28.4x19.8x2.5 cm).

The only catch: its six-cell battery juts out at the back, whereas some netbooks, such as the Eee PC 1005HA, have comparable battery sizes that fit more flush with the system.

The keyboard

The Aspire One 751h has raised keys with a slightly textured feel. The layout is close to full size, and we appreciate that the right *Shift* key is of full size and in the proper place. Touch typists will have no problem with this netbook, but we noticed that the keyboard flexed as we pounded out responses to IMs.

The touchpad and touch button

When we had reviewed the latest 10-inch Aspire One, we had complained about the tiny touchpad. With the 751h's slightly wider footprint, however, comes a wider touchpad (2.5 x 1.6 inches or 6.3 cm x 4.1 cm), which was a pleasant improvement over the D250's 2.0 x 1.5-inch (5.1 cm x 3.8 cm) pad.

Display and sound

One feature the 751h offers that most of its 10-inch competitors don't is a high-resolution screen. Whereas the 1005HA and Aspire One D250 have displays with 1024 x

600-pixel resolution, the AO751h's 11.6-inch display has a resolution of 1366 x 768. Those extra vertical pixels, in particular, mean you won't have to scroll down as often when you're viewing pages, a common inconvenience with netbooks.

The volume, as you would expect with a netbook, is weak. Music was predictably tinny, but no worse than on other netbooks.

Ports and the webcam

The 751h has a standard selection of ports: three USB ports, VGA output, an Ethernet jack, headphone and mic ports, and a 4-in-1 memory card reader. The memory card reader, VGA and one USB port are on the right side, and the rest are on the left.

The netbook comes with an embedded 1.3 mega-pixel webcam, is exceptionally bright and sports extremely high latency.

Processor

The AO751h has most of the components that a netbook normally features—1GB of RAM, a 160 GB hard drive, and either a 1.6-GHz or 1.66-GHz Intel Atom processor. However, this one sports a 1.33 GHz Intel Atom Z520 processor. Like many other netbooks, the 751h uses Intel's integrated GMA 500 graphics solution, with 384 MB of shared memory.

Battery life

The 751h has some tough competition in the battery life department. I found the 1005HA lasting nearly four hours on JoliCloud. The 751h lasted three and a half hours, which, given that Linux has absolutely messed up power management capabilities, is still an impressive score. It beats my official six-cell HP notebook's average by almost 90 minutes.

Verdict

If you're craving a little more screen real estate than the average netbook, along with a high-resolution display, the Acer Aspire One 751h combines several compelling features, such as hours of battery life and an exceptionally slim 1-inch silhouette. But if all you want is a netbook—and not necessarily an 11-inch one—there are several comparably priced (albeit, thicker) 10-inch models, such as the Dell Mini 10V and the ASUS 1005HA, which run faster and last even longer.



ASUS Eee PC 1005HA

If I'm not mistaken, ASUS must have launched more than 20 models of its popular Eee PC netbook in the past year-and-a-half. With that kind of track record, it's no surprise that this latest entry is one of the most well-rounded netbooks on the market. The ASUS Eee PC 1005HA has a chassis similar to the elegant Eee PC 1008HA and the same Intel Atom processor, 1 GB of RAM, and 160 GB hard drive as the 1008HA. And its high-capacity battery lasts just short of four hours on JoliCloud, making this one of the longest-running netbooks yet.

Inspired design

A member of ASUS' Seashell family, the 1005HA is the 1008HA's chunkier twin. Though less svelte than the inch-thick 1008HA, the 1005HA draws from the same gene pool—a fact demonstrated by its curved panels and sharp look. Gone, however, is the uniform body of the 1008HA; no plastic covers the edges or hides the ports. Its glossy lid is covered in a scratch-resistant in-fusion finish, which is far from fingerprint-resistant, so you'll need to keep a cloth handy.

The vital stats measure 10.2 x 6.9 x 1.4 inches (25.9 x 17.5 x 3.5 cm). It also weighs noticeably more than the 1008HA, but the 1.4 kg 1005HA netbook didn't strain my shoulders even with the charger.

Ports

On the left side are a USB and full VGA port. On the right side are a 3-in-1 card reader, an Ethernet port, two USB ports (bringing the USB total to three, one more than the 1008HA), and headphone and microphone jacks.



The keyboard and touchpad

The Eee PC 1005HA includes the same comfortable, spacious keyboard found on the 1008HA. However, there are no spaces between the raised black matte keys, which I found flat (I prefer them raised), but they did provide nice, tactile feedback. Moreover, the right *Shift* key is full sized and directly below the *Enter* key. A touchpad on/off button sits at the top left of the keyboard.

The 2.5 x 1.5-inch (6.3 x 3.8 cm) touchpad blends

in with the system's deck and is covered with small braille-like dots. Unlike most netbooks, the pad had very little friction, and navigating the desktop didn't require much backtracking. As with other Eee PCs, the touchpad also supports multi-touch gestures; using two fingers to pinch the pad and zoom in on a Web page was easy, and the image quickly rendered in its new size. The single mouse button was comfortable for making selections.

Display, audio, and the webcam

The 10.1-inch (25.6 cm), 1024 x 600-pixel resolution LED-backlit display on the 1005HA allows Web pages and windows to fit-to-size on the screen. While the videos looked smooth, and the colours were bright, tilting the glossy screen back 45 degrees resulted in a slight glare; horizontal viewing angles were better. From afar, the inch-thick, glossy black bezel looked as if the screen was flush with the frame. Up close, however, you could clearly make out the raised frame.

Above the display, a 1.3 megapixel webcam provided clear images. The speakers, located at the bottom front edge of the system, allowed us to hear the soundtrack loud and clear.

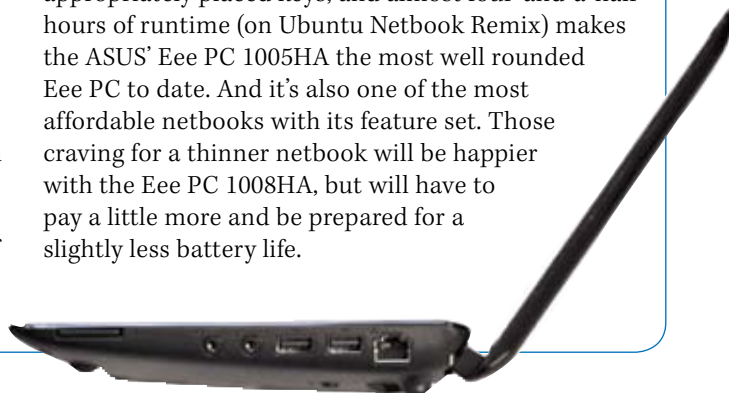

The processor

Configured with the same 1.66-GHz Intel Atom N280 processor and 1GB of RAM as the ASUS Eee PC 1008HA, I didn't see much change in everyday performance on the 1005HA.

The Intel GMA 950 integrated graphics chip with 128 MB of shared memory was able to playback a downloaded high-definition 720p video clip smoothly with no hiccups.

Verdict

An attractive design, comfortable keyboard with appropriately placed keys, and almost four-and-a-half hours of runtime (on Ubuntu Netbook Remix) makes the ASUS' Eee PC 1005HA the most well rounded Eee PC to date. And it's also one of the most affordable netbooks with its feature set. Those craving for a thinner netbook will be happier with the Eee PC 1008HA, but will have to pay a little more and be prepared for a slightly less battery life.



ASUS Eee PC1008HA

The most amazing thing about the Seashell 1008HA is the design, which is based on its name—rounded and clamshell-like. Measuring just 27mm at its broadest point, and getting considerably thinner towards the front of the chassis, you'll hardly notice the 1008HA Seashell in your bag, which is also because it weighs just over 1 kg.

Display

The 10-inch (25.4 cm) screen is a joy to use. It features an average resolution of 1024 x 600 pixels, but images appear perfectly crisp, and the colour reproduction is flawless. That said, the glossy Super-TFT coating can make it difficult to view in the brightest conditions, with reflections proving problematic.

Design

Despite the Eee PC 1008HA's compact footprint, it features a well-spaced keyboard that stretches right to the edge of the chassis. The keys are 92 per cent full-sized, and feature a flat design with slightly tapered edges. It's not the best netbook keyboard we've seen, but it works well, and is very tough. It's also a comfortable option when typing for longer periods.

A high quality standard runs throughout, although the high-gloss plastics do attract fingerprints more than rival products, and you'll have to use the protective case

provided if you want to keep it looking new. The svelte lines are enhanced by hidden ports, protected by plastic covers. This feature doesn't help usability, however, with access proving slightly tricky.



▲ All the ports are stylishly covered to prevent dust and to give it a finished look

Ports

Behind the covers you'll find a USB port on each side of the chassis, and on the left sits a mini-VGA port. It's the first time we've seen one of these in action, but thankfully an adapter is provided as standard, slotted neatly underneath

the chassis. On the other side, you'll also find microphone and headphone sockets, and an Ethernet port.



▲ The lan port is placed too far away at the corner to make it a bit inconvenient

Storage

The ASUS also impresses when it comes to storage space, with a capacious 160 GB hard drive included as standard. You'll find an additional 10 GB of online storage provided—offering better protection for your most important files, with no risk of damage if anything happens to the netbook itself.


The processor

Processing power comes from an Intel 1.66 GHz Atom processor, backed by 1024 MB of memory. Performance is reasonable, as long as you don't intend to carry out intensive multi-tasking. Mobility is excellent, with the compact 3-cell battery life offering almost four between charges under normal workload.

Verdict

Overall, the Eee PC Seashell 1008HA is a compelling proposition. Offering quality, comfort and excellent mobility, it's one of the best mid-range netbooks you can currently buy.

Comparison between 1005HA and 1008HA ►



THE COMPLETE MAGAZINE ON OPEN SOURCE





LINUX BETA
ForYou

Your favourite Linux Magazine is
now on the Web, too.





LinuxForU.com

Follow us on Twitter @LinuxForYou

Comparative spec sheet for popular netbooks

BRAND	Acer Aspire One		ASUS EeePC	
				
MODEL	D250	751H	1005HA	1008HA
Processor	Atom N280 (1.68GHz, 667 MHz FSB, 512KB L2 cache)	Atom Z520 supporting Intel 32 Architecture	Atom N270 / N280	Atom N280
Chipset	Intel 945GSE	Intel US15W		
System Memory	1 GB DDR2 (soDIMM)	1 GB DDR2 (soDIMM)	1 GB DDR2 (soDIMM)	1 GB DDR2 (soDIMM)
Display	10.1" LED TFT (WSVGA, 1024x600)	11.6" LED TFT (WXGA, 1366x768)	10.1" LED (WSVGA, 1024x600)	10.1" LED (WSVGA, 1024x600)
Storage	160GB 2.5" SATA II 5400RPM HDD	160GB 2.5" SATA II 5400RPM HDD	160GB 2.5" SATA II 5400RPM HDD	160GB 2.5" SATA II 5400RPM HDD
Audio	High-definition audio support; Two built-in stereo speakers; Built-in digital microphone	Dolby Headphone; High-definition audio support; Two built-in stereo speakers; Built-in digital microphone	Hi-Definition Audio Codec; Stereo speaker; Digital Array Mic	Hi-Definition Audio Codec; Stereo speaker; Digital Array Mic
Camera	0.3 megapixel	1.3 megapixel	1.3 megapixel	1.3 megapixel
Wireless	WLAN: 802.11b/g; ETHERNET: 10/100Mbps; BLUETOOTH 2.0 + Enhanced Data Rate (EDR)	WLAN: 802.11b/g; ETHERNET: 10/100Mbps; BLUETOOTH 2.0 + EDR	WLAN: 802.11b/g/n; BLUETOOTH 2.0 + EDR	WLAN: 802.11b/g/n; BLUETOOTH 2.1 + EDR
Dimensions	258.5mm (W) x 184mm (D) x 25.4mm (H)	284mm (W) x 198mm (D) x 25.4mm (H)	262mm (W) x 178mm (D) x 25.9mm ~ 36.5mm (H)	262mm (W) x 178mm (D) x 18mm ~ 25.7mm (H)
Weight	1.27 kg (with 6-cell battery pack)	1.35 kg	1.27 kg (with 6-Cell battery pack)	1.1 kg
Power	4hours, 6-cell 57.7Wh Li-ion battery	4.5hours, 6-cell 57.7Wh Li-ion battery	5hours, 6-cell 63W Li-ion battery	3.5hours, 6-cell 32Wh Li-ion battery
I/O Ports	Three USB 2.0 ports; External display (VGA) port; Headphone/speaker/line-out jack; Microphone-in jack; Ethernet (RJ-45) port; DC-in jack for AC adapter	Three USB 2.0 ports; External display (VGA) port; Headphone/speaker/line-out jack; Microphone-in jack; Ethernet (RJ-45) port; DC-in jack for AC adapter	1 x VGA Port (D-sub 15-pin for external monitor); 3 x USB 2.0; 1 x LAN RJ-45; 2 x Audio Jack (Head Phone / Mic-in); Card Reader: MMC/SD(SDHC)	1 x Mini VGA Connector; 2 x USB 2.0; 1 x LAN RJ-45; 2 x Audio Jack (Head Phone / Mic-in); Card Reader: MMC/SD(SDHC)
Price	Rs 20,200	Rs 24,000	Rs 24,500	Rs 28,000

Comparative spec sheet for popular netbooks

HP	Dell	Samsung		Sony Vaio	BRAND
					
Mini 110	Inspiron Mini 10V (with Ubuntu 8.04)	N310	N140	W Series	MODEL
Atom N270 / N280	Atom N270 / N280	Atom N270	Atom N280	Atom N280	Processor
Intel GMA 950	Intel GMA 945GSE		Intel GMA 950	Intel GMA 950	Chipset
1 GB DDR2 (soDIMM)	2 GB DDR2 (soDIMM)	1 GB DDR2 (soDIMM); 128MB shared graphics memory	1 GB DDR2 (soDIMM)	1 GB DDR2 (soDIMM)	System Memory
10.1" LED (WSVGA, 1024x600) / 10.1" HD LED (1366x768)	10.1" LED (WSVGA, 1024x600)	10.1" LED (WSVGA, 1024x600)	10.1" LED (WSVGA, 1024x600)	10.1" LED (WXGA, 1366x768)	Display
160GB/250GB 2.5" SATA II 5400RPM HDD; 32GB (Solid State Drive Flash Module)	120GB / 160GB 2.5" SATA 5400RPM HDD; 16GB (Solid State Drive Flash Module)	160 GB 2.5" HDD 5,400 RPM	250GB 5400 RPM	250GB Hard Disk Drive	Storage
Built-in Stereo speakers	Main Speakers (quantity) : 2 x 1.0W	HD Audio, Mic Noise Suppression, SRS 3D Sound Effect	HD Audio, SRS TruSurround XT, SRS WOW XT, SRS CS Headphone	Intel High Definition Audio Sound System	Audio
1.3 megapixel	1.3 megapixel	1.3 megapixel	0.3 megapixel	1.3 megapixel	Camera
WLAN: 802.11b/g/n; BLUETOOTH 2.0 + EDR	WLAN: 802.11g/n; BLUETOOTH 2.1 + EDR	802.11 b/g; Bluetooth 2.1 + EDR	WLAN: 802.11b/g/n	WLAN: 802.11b/g/n; Integrated Stereo A2DP	Wireless
10.3in (W) x 6.77in (D) x 1.04-1.29in (H)	261mm (W) x 182mm (D) x 26.8mm (H)	10.31in (W) x 7.26in (D) x 1.1in (H)	10.3in (W) x 7.3in (D) x 1.01in (H)	10.54in (W) x 7.07in (D) x 2.01in (H)	Dimensions
1.1 kg	1.2 kg (with 6-cell battery pack)	1.5 kg	1.45 kg	1.6 kg	Weight
3 / 6 Cell Lithium Ion Battery	6-cell 56Wh Li-ion battery	6-cell 56Wh Li-ion battery	6-cell 52Wh Li-ion battery	3.5 hrs; 6 Cell	Power
3 USB Ports, Integrated 10/100BASE-T Ethernet LAN (RJ-45 connector), 5-in-1 integrated Digital Media Reader	3 USB Ports, Integrated 10/100 Ethernet LAN (RJ-45 connector), 3-in-1 Card Reader	3 USB Ports, VGA; 3 in 1 (SD / SDHC / MMC)	3 USB Ports; VGA; Integrated 10/100 Ethernet LAN (RJ-45 connector); 3 in 1 (SD / SDHC / MMC)	2 USB Ports; VGA;	I/O Ports
Rs 21,499	Rs 21,840	Rs 22,990	Rs 23, 499	Rs 28,000	Price

Moblin 2.1: Linux for your mobile devices

Moblin is short for 'mobile Linux'. Well, it's also a Linux-based OS (the official release is a Fedora derivative) that's optimised for small computing devices. Naturally, this calls for a complete UI overhaul—nothing like the GNOME and KDE we're used to and fight over.

The Linux Foundation released Moblin 2.1 on November 5, which made us take it for a test drive on the netbooks.

Installation

The installation procedure that worked for us:

1. Download the image file from <http://moblin.org/downloads>.
2. Insert a USB drive (minimum 1GB)—make sure the drive is unmounted—and run the following command as root:

```
dd bs=4096 if=/path/to/image-file of=/dev/sdX
```

...where X in sdX is the number for your connected USB device.

3. Now, you can boot your machine off the USB drive and install Moblin just the way you would, for any other distro. Yes, it's that simple!

Faster boot... and shut down?

On the ASUS Eee PC 1008HA (as it is configured today as I write this), from the time the power button is pressed until the machine is up and ready to go is 30 seconds. BIOS consumes 8 seconds, and Moblin boot consumes an additional 22 seconds. The shut down time for the OS was 8 seconds. This runs in sharp contrast to the boot times reported by the guys at Moblin, on their internal builds, which are already in the sub-10 second boot range.

Users want much faster boot times from cold starts. This is something Intel (and Google as well) recognise, and is a major area of focus today.

Initial impressions

The team working on Moblin has not been complacent in their duty to provide us with a powerful new GUI,

plenty of attractive eye-candy and fast operating graphical features. Even so, and to address the graphical needs of a full-on Moblin experience, the team realised that more software support was needed than was possible, at least not without digging deep into the innards of Linux and seeing what's there and what needed to be changed—and that's exactly what they did.

The traditional desktop environment is replaced in Moblin. Yet, the desktop should seem to be highly GNOME based. At the top of the screen is a bar which appears either when you press the Super key, or move the mouse towards the top. It auto-hides graphically, which can sometimes be very annoying since once one gets used to the environment, having this graphical sequence of sliding down 3/4-inch, even though it takes only about half a second, is not amusing.

A completely revamped UI

Graphically, Moblin is very well designed. It's attractive, and uses the available space quite well. Everything has a good default contrast ratio, and colours appear vibrant on both icons and dialogues. It's very clear the Moblin design team wanted a powerful visual experience. Honestly though, I think the team went a little overboard with some of the GUI effects.

For example, every time the mouse is moved over one of the icons, it bounces as the tool tip text appears beneath it. This is extremely annoying to me as a consistent visual trait that appears, no matter how fast I mouse over something and click on it. I'm sure each graphical effect takes only about 1/8 second or so, but after the initial sheen is gone, the delay and bouncing graphics starts to become an eye sore. A nice feature would be the ability to turn this graphical trait off.

Alt+Tab cycles through open applications in a pleasing graphical way, with a thumbnail of each being shown. In many cases, this is the only way to get to the other applications that are running since only one



▲ MyZone—the default screen



▲ The Gecko-based Web browser



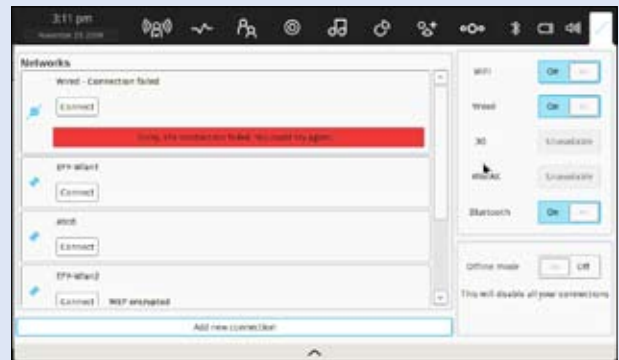
▲ Application Manager

application can be visible on-screen, at one time.

I guess, it's now time to concentrate on the buttons sported by the top bar.

MyZone

In detail, this area shows the latest applications which have been run, allowing for single-click, or eventual single-touch access. It has a smart application area that features the calendar, to-do list and links to favourite applications. There is a column on the right side which shows not only thumbnails of recently used programs, but also thumbnails of news items and other feeds which can be fetched off social websites. The middle column showcases thumbnails



▲ The half-baked network manager called ConnMan (the name is surely an irony)


for the websites frequently visited (and media files recently played) by the user.

Status

With a few Web services accounts, you can update your online status without actually visiting the sites. As of now, services like Google chat and Twitter have been incorporated.

People

This tab is an interface to all your online contacts. Of course, it currently uses your your accounts configured in Empathy to assimilate the contact list.



India's Premier Linux & OSS Academy
in Associate with

EnterpriseDB
The Enterprise Postgres Company



The Official
EnterpriseDB

PostgreSQL Training

Duration : 3 Days - Dates: 18, 19, 20 Dec 2009

Register Now : Limited Seats

Email : training@enterprisedb.com / contactus@lynusacademy.com

044-42171278
9952966527

contactus@lynusacademy.com
www.lynusacademy.com



Lynus Academy Pvt Ltd

1961-B, Vijaya Complex - 5th Floor
Asiad Colony, Anna Nagar - West
Chennai 600101 INDIA

Introduction Offer for Online Training

- ♦ Shell Scripting - 16 Hrs - Rs. 2500 /-
- ♦ PHP+MySQL - 40 Hrs - Rs. 6500 /-
- ♦ ROR - 40 Hrs - Rs. 6500 /-
- ♦ Drupal / Joomla - 16 Hrs - Rs 4500 /-

Limited Time Offer ..

Call for Demo Now ..

Internet

This tab features your most favourite websites and also an address bar that launches the Web browser. Speaking of the Web browser, it is a modified version of the popular Gecko-based Mozilla Firefox. After the launch of Moblin v2.1, the browser has started to exhibit much better performance, along with support for plug-ins and add-ons.

However, there is a glitch here. I am a bit accustomed to the idea of typing away right after I press Ctrl+t to open a new tab. Here, the new tab shows *moblin://start/* on the address bar, with cursor after the slash. The implication of the same is, I need to physically select the text and then type, which I found to be pretty irritating. Interestingly, this was not the case in the Moblin we used on Mandriva (install the *task-moblin* package on Mandriva 2010 and you're done). There, *moblin://start/* was selected by default.

Media

The media player has a long way to go before it's ready for prime time. However, the underlying concept of accessing media content is quite brilliant.

An initial launching screen is presented, which allows the single-click launch of media content which is on the screen. This screen remembers your last search settings and presents things exactly as they were when you last closed the app. A search bar is presented allowing rapid access to the content you're looking for.

Accessing images automatically comes in a frame-by-frame viewing format, or they can be displayed in a continuous slide show. Videos are played from beginning to end after being selected, unless cancelled. However, if you let it finish then a bug in the current media player restarts playing content at the beginning of all media, rather than the next item in your search from where you launched the audio or video clip.

The media player interface hides all controls during viewing, and presents an automatic full screen experience. However, there are recurring bugs throughout, which prevent content from being shown when clicked, only heard, the pause doesn't always work properly, etc.

Pasteboard

Pasteboard is just another clipboard manager. However, something beautiful about this one is, it is persistent. That is, the contents of the pasteboard sustain across reboots.

Applications

The application tab houses a list of all the applications that are stored in your OS, along with links to Moblin's settings utilities.

Zones

This section helps you to manage, organise and switch between different applications. While I am still clueless about the usefulness of this 'feature', it just might be Moblin's attempt in emulating the concept of multiple desktops.

System tray

The system tray houses the Bluetooth settings (not available on previous versions), power options, screen brightness and connectivity information.

Finally...

At the outset, I am pretty pleased with Moblin. It is still very shaky as there are errors which pop up during normal use, at least once every hour. Even so, I can see the amazing and tremendous progress and potential in every aspect that I've seen so far.

The Moblin.org team has gone to great lengths to rethink the very concept of a user interface, even addressing future technologies like netbook touch-screens, and their effort shows in great regard. Moreover, the portability of netbook devices, their long battery life and their feature set make it a perfect target for a user-friendly Linux-based operating system.

My love and hate relationship with Moblin

1. No application is running, the MyZone screen should re-appear automatically, rather than forcing the user to press the Super key or moving the mouse to the top. After all, if a user is using his PC, what good is it to stare blindly at the desktop background image after closing an application?
2. The battery widget doesn't indicate the battery life in units of time. A mere percentage is hardly informative.
3. The nomenclature for the sound and brightness levels are uber cool.
4. The network manager (it's an application called ConnMan) is a major pain in you-know-where. While it has no issues using DHCP, there is absolutely no way to set up a static IP. Of course, you can always move to the console, but even there, we failed to associate the ESSID of the wireless network.
5. The new version of Moblin finally supports resolutions higher than the standard 1024x764.
6. There is no 'shut down', 'hibernate' or 'sleep' button on the UI. Only thing you can do is shut it off, by pressing your netbook's physical power button.
7. The only power save option available is to set the duration of inactivity before it goes to sleep.
8. Application menu's UI is innovative but not too handy. Thankfully, pressing Alt+F2 takes you to the application menu's search box.

Hop onto the (Joli)Cloud

Jolicloud transforms your netbook into a sophisticated Web device that taps into the cloud to expand your computing possibilities. The Web already hosts a significant part of our lives: e-mails, photos, videos and friends are already somewhere online. Jolicloud was built to amalgamate the computer and Web part of the same experience.

The OS was created by Tariq Krim, founder and ex-CEO of Netvibes. The developer team has a huge inclination towards the Web, so they built the user interface mostly using its core technologies (HTML, JavaScript and CSS). They have integrated our Web DNA into the OS to make it modular, social and personal. The developer platform relies on the Web and will let anyone or any service join in no time. With their API, developers will have the ability to let their website communicate with the computer directly, with no need for code specific native applications.

How to get it

Unfortunately, as of writing this article, JoliCloud is in private alpha. However, you can ask for an invitation link at <http://url.shayonpal.com/jolicloud>.

After you get the invitation code in your inbox, use that to get registered (keep the credentials, you will need them later) and download the ISO file. After that, download the USB Creator software from <http://url.shayonpal.com/JoliUSBCreator>.

Installation

Now, fire up the terminal and go to the directory where you downloaded the JoliCloud USB Creator.

```
cd /the/path/to/your/directory
chmod a+x jolicloud-usb-creator-1.1.2.sh
./jolicloud-usb-creator-1.1.2.sh
```



▲ JoliCloud USB Creator

At the last step, you will be asked for the admin password.

1. Now, connect the thumb drive to your computer.
2. In the JoliCloud USB Creator window, click *Browse* to find and select the Jolicloud image you've downloaded.
3. Select your target device: The USB key you have inserted should be automatically detected.
4. Click *Create* to create your Jolicloud USB.

We have finally managed to create a bootable USB drive with a copy of JoliCloud in it. Now, insert the USB key into your computer and boot off it. In the boot menu, select 'Install JoliCloud'.

Activate your JoliCloud account

Yes, this is what will distinguish JoliCloud from any other operating system. Do you remember that you had registered on the JoliCloud website? Well, we will need those credentials now.

Get your netbook connected to the Internet and then click on the little cloud on the system tray. Alternatively, you can also click on 'Get Started' on the 'Favourites' launcher menu.

Enter your username and password, and you are connected with the (Joli)Cloud.

The desktop

Upon booting Jolicloud, you'll notice instantly that



▲ The default/home screen



▲ Logging into the JoliCloud's Web service.



▲ The application manager window



▲ Updates from your 'connections'

it's basically Ubuntu with some modifications. Specifically, it's based on Ubuntu Netbook Remix (UNR) version 9.04. In fact, the default login sound, which seems quite out of place in the modern-looking Jolicloud, hasn't even been changed.

Jolicloud keeps the launcher application and basic design of UNR v9.04, too. If you're not familiar with the older UNR, this basically means three things:

- First, all your applications are sorted into category-based tabs. Clicking on a tab displays a page with large icons for each application. A single click opens the application. Additionally, on the right of the screen are links to your home folder, documents folder, and so on. These folder links are not as well integrated, though, and open up an instance of Nautilus.
- Second, all windows are maximised (or should be). Rather than letting users place applications side by side, UNR forces users to work with one application at a time. This actually makes a lot of sense with some applications, because there really isn't enough space to work with multiple applications on a netbook. Other applications, though, don't work as well this way.
- Finally, the windows list only displays the full name of the current application. The panel at the top of your screen puts all the instances you're

running neatly into tabs of icons, making for seamless switching between applications.

A deeper look into the OS

JoliCloud uses Mozilla's Prism Web runtime and UNR to deliver a Web-centric Linux environment that is easy to use. The platform and its associated Web service are at an early stage of development and as mentioned earlier, are currently in closed alpha testing.

JoliCloud introduces new features with a new interface for installing applications. For starters, the developers have had the sense not to offer applications like GIMP, which really doesn't work well with the concept of 'one window at a time'. More importantly though, JoliCloud makes installing applications a one-click process. After clicking the install button under any application, it gets added to a queue of software updates and installs. This means no more waiting for Synaptic to close before performing updates.

Even cooler, though, is the integration of Web applications into the application area. There is no distinction between standard applications and Web applications (which run in Mozilla Prism). When you are offline, these applications will still launch, but can't do much, unless they use Google Gears to store data offline.

However, the extensive use of Prism largely defines the philosophy behind the JoliCloud user experience. Web applications are treated like conventional software—each one is launched in its own window and process. This is perhaps similar to the model that Google will pursue with its upcoming Chrome OS platform.

The only significant complaint I have with JoliCloud is that there are some rough spots in the theme. It looks great in all the screenshots on the JoliCloud website, but there are some hidden (and not so hidden) unthemed parts.

For example, the volume control is taken straight out of default GNOME with no theming. It integrates just fine with Ubuntu, but it doesn't look right with the dark, modern look of JoliCloud.

JoliCloud centres on a directory of applications that can be sorted by genre, release date, and popularity. To download or remove them from your computer, just click on their icon and it does the rest. Jolicloud groups both Web apps and software programs under the same name umbrella, and both are added and removed from your system in the same manner. There're also normal Add and Remove programs tools just like you get in Windows, but it's easier to use them from JoliCloud's rounded and simplistic interface.

What I liked

- Absolutely hassle-free and all the hardware functioned out-of-the-box.
- Some bits and pieces are still in beta, but the underlying approach of making websites and software applications feel the same, as well as introducing users to new ones, is really innovative.
- Has the heart of Debian.
- Can easily be converted into a bootable USB.
- Multiple styles of application installation.

What I am skeptical about

- It relies too heavily on Ubuntu's netbook remix interface, to the extent that it still seems to be Ubuntu. I'd like to see the overall functionality integrate further before I'd call this its own markedly different initiative.
- Downloaded applications are organised for you and cannot be reorganised.
- Lousy power management.
- If you add the JoliCloud APT repository to the sources list in a regular Ubuntu 9.04 desktop install, you can easily install packages for various Prism launchers without even having to bother installing the JoliCloud distro or using its Web service. This gives you roughly the same experience minus the slightly creepy social networking features. So, what's the big deal?

you can ask friends who have the same stuff or the same machine.”

This should provide a great way to get to know the best applications when the list of tools eventually grows too big for users to easily find their way. It's also a good way for JoliCloud to analyse who is using what, and on what type of machine (this reminds me a bit of Wakoopa). It could prove beneficial for JoliCloud to integrate more social features in the future, like the ability to communicate directly with other users or instantly share applications, reviews, screenshots, etc, on the multitude of social networking services out there.

One of the big draws of

JoliCloud's launcher and user interface are entirely based on UNR, but it has a few minor theming customisations that differentiate it aesthetically. The most significant feature of JoliCloud is its unique Web service that runs in a Prism shell and integrates with the rest of the environment. It provides an elegant HTML-based interface for update management, an application directory, and some simplistic social networking features.

The JoliCloud service's Dashboard interface shows you notifications from the service and information about pending updates. Users can initiate a system update from the JoliCloud dashboard interface, but the actual update process is handled by Ubuntu's standard package update tool.

An OS gone social

This is the most intriguing part of JoliCloud that differentiates it from the rest of the operating systems. While the JoliCloud servers keep the data about the software that is being installed by you, it is also possible to connect with various other JoliCloud users and stay updated on the kind of software they are using. Yes, this is both exciting and scary at the same time. While privacy could be a serious concern among users, it is also a great way to get introduced to new apps and services.

As Tariq Krim said, “You are right. We are dropping the netbook remix for a newer interface and we are going to improve the social features. We also see the social OS as a great way to do support as

JoliCloud is that it takes this list of apps you have installed and backs it up. If you have multiple computers running JoliCloud that share the same account, it syncs up those apps, including any log-ins or shared data. This places less importance on what hardware you're using, so you can hop from machine to machine and get right back to what you were doing, on any one of the systems.

Bundled software

Some of the software that come bundled with JoliCloud are:

1. Mozilla Firefox
2. Evolution
3. Cheese
4. Empathy
5. F-Spot
6. Transmission
7. Rhythmbox

In short, almost anything and everything that a netbook might need, is pre-installed with the OS.

Compatibility issues

True to the Ubuntu charm, JoliCloud worked perfectly with every netbook we tested it on. All the drivers and devices got installed and could be used out-of-the-box. However, we could not make it presentable on the 11.6” (29.5 cm) Acer Aspire 751H because the resolution could not be changed from 1024x768 and the icons looked extremely stretched. Other than that, it was the most hassle-free OS we tested.

Remixed Ubuntu for the netbooks

Ubuntu Netbook Remix joins Moblin as a lightweight, optimised solution for Atom-based netbooks, while offering the usual stability and friendliness of the Ubuntu family.

In this short review, we'll look at what Ubuntu Netbook Remix (UNR) has to offer. From wireless to multimedia, it's Ubuntu all over. The major difference is in the desktop layout and the performance on low-end modern processors.

Get UNR

UNR is available for download from the official Ubuntu website. However, getting UNR ready for use is slightly different from the typical Linux distribution, so let's spare a few moments and see how you can get UNR booted on your machine.

After you have downloaded the UNR v9.10 ISO file, you will have to copy it, bit-by-bit to a bootable external device. Standard USB thumb drives are ideal for this task. Images can be copied using the `dd` command—or much simpler, by using the USB Starter Disc Creator utility, which can be found in the Ubuntu Live CD. Follow the instructions and you're all set.

The desktop

The UNR desktop has a very unique, interesting look. The choice of colours is Ubuntu, with browns, oranges, and a darker, more sombre scheme than the original. What makes the biggest difference is the layout, which takes some time getting used to. On the left side, there's the GNOME menu, with combined *Applications*, *Places* and *System* menu combined. The right side displays the contents of the menu.

As expected, by default, the screen focuses on the *Favourites* menu. It comes with a handy set of default favourite apps. Of course, it's up to the user to customise it to his own likes. You simply hover the mouse over an application icon and click the plus sign that appears to define it as one of your 'favourites'. In fact, during my tests I was able to include all the apps I required here so that I hardly required to navigate down the other menus.

The top panel remains with the standard system tray, populated with icons for wireless connectivity, Bluetooth, battery levels, etc. However, what seems like the task bar is actually a place for the title bar of the application you're currently working on. This is a good add-on which saves some essential screen space. But, how to access the other apps that are also open? Well, they're listed as small icons on the left the



▲ The default/favorites tab on UNR



▲ Accessing the Games menu



▲ Running Firefox; note the title bar fuse into the panel



▲ The background gets dimmed when an unmaximised window is open

window manager part of the panel. The Ubuntu icon on the extreme left serves as the 'Show Desktop' utility.

Applications

Clicking on the categories in the left column, you get the standard Ubuntu collection. You have the OpenOffice.org suite, Totem, Rhythmbox, the GIMP, and other programs you're used to, from the stock release. Likewise, expanding the *System* section, you have everything you would expect from the Ubuntu. As always, installing new programs is dead easy with Synaptic.

You won't get codecs for proprietary formats out-of-the-box, but you will be able to download and install them the first time you try to play a file in one of these formats, like mp3, for instance. Similarly, Flash is not available by default, but you can install it either manually or via the Medibuntu repository.

To sum up

Overall, UNR is a solid product. It's Ubuntu, dedicated to the 'low-specs' market with very handy UI enhancements that make it ideal for devices with lesser screen real estate. If you have a netbook and don't particularly care for a sub-optimal Windows XP installation, nor do you wish to dabble with the brand-new Moblin 2.1, then Ubuntu Netbook Remix seems like a good idea. I believe you will like the result. The

combination of open source, lots of good programs, stability, usability, and decent looks is really hard to beat.

Of course, there's a Kubuntu Netbook Remix too for the hard-code KDE4 users. **END** 



Kubuntu Netbook Remix 9.10. Read the review online @ www.linuxforu.com/reviews/kubuntu-netbook-remix-review

By: Sayantan Pal

An avid Twitter user and a social media enthusiast, the author is a passionate blogger and a professional gamer too. He also feels compelled to be opinionated about anything that comes his way, be it Linux distributions, our marketing strategies, table etiquettes or even the fabled Ramsay movies!

Surveillance & Security ZONE

Securing your Business



South Asia's largest ICT event

18th
Convergence India 2010

International Exhibition & Conference
Pragati Maidan, New Delhi, India
23-25 March 2010

Supports



Department of Telecommunications
Ministry of Communications & Information Technology
Government of India



Department of Information Technology
Ministry of Communications & Information Technology
Government of India

Organiser

Ei Exhibitions India Pvt. Ltd.
ISO 9001:2008 & ISO 14001:2004

Vipul Mehra, General Manager, vipulm@eigroup.in
Tel: +91 11 4279 5000 / 5030 M: +91 99102 04131

Supporting Journal

LINUX
LIFE

Certified by



ufi
Approved
Event

IAEE
International Association
of Exhibition Events

IEIA
Indian Exhibition
Industry Association

www.convergenceindia.org



CODECHEF

Welcome to the CodeChef Challenge—your monthly dose of coding puzzles from India's biggest online programming contest, now in print!

Well, it's time to again compete with like-minded folks who are as passionate about programming as you are. Each month, we will bring to you a CodeChef challenge that will test your tech skills, and the top entries will be awarded by *LINUX For You* and Directi.

Last month's puzzle

How many squares of different sizes, and with sides parallel to the sides of the chess board, can we get from a 8x8 chess board?

The solution

An 8x8 chess board has 204 squares of different sizes with sides parallel to the sides of the chess board.

How we got the answer?

Answer is equal to $F(1) + F(2) + F(3) + \dots + F(8)$, where $F(n)$ is the number of squares with side 'n' units on the chess board.

Consider the top left corner of each square. The position of the top left corner along with 'n' can completely define a square. If the side of the square to be placed is 'n' units, then we can place the top left corner in $(8 - n + 1)$ ways on any row. So, for 'n' = 8, we can place the top left corner in $(8 - 8 + 1)$, i.e. 1 way on any row.

The same is true for the columns.

So, for a 8 x 8 square, we can use only the first row and column, for 7 x 7 squares, we cannot place the top left corner on any row or column number beyond 2.

Thus, the total number of squares of side 'n' units is $(8 - n + 1) \times (8 - n + 1)$. This is also the value for $F(n)$. Using this formula, the values of $F(n)$ for different values of 'n' are:

$$F(1) = (8 - 1 + 1) \times (8 - 1 + 1) = 64$$

$$F(2) = (8 - 2 + 1) \times (8 - 2 + 1) = 49$$

$$F(3) = (8 - 3 + 1) \times (8 - 3 + 1) = 36$$

$$F(4) = (8 - 4 + 1) \times (8 - 4 + 1) = 25$$

$$F(5) = (8 - 5 + 1) \times (8 - 5 + 1) = 16$$

$$F(6) = (8 - 6 + 1) \times (8 - 6 + 1) = 9$$

$$F(7) = (8 - 7 + 1) \times (8 - 7 + 1) = 4$$

$$F(8) = (8 - 8 + 1) \times (8 - 8 + 1) = 1$$

Add up all the results and you get 204


And the winners are...

- Wallace Jacob
- Anuvrat Parashar
- Arjun Pakrashi

Puzzle of the month

How many numbers from 1 to 10000 (inclusive) are divisible by 3 or 5 or 7?

Want to win?

Answer the question of the month for the CodeChef Challenge and mail it to codechef@efyindia.com. Three lucky winners will win Rs 1,000 each in cash! 

About CodeChef: CodeChef.com is India's first, non-commercial, online programming competition, featuring monthly contests in more than 35 different programming languages. CodeChef has been created by Directi, as a way to continuously challenge and engage the developer community. The goal is to provide a platform for practice, competition and improvement as well as enable developers to benchmark their skills against their peers. Log on to www.codechef.com for the December Challenge that takes place from December 1-11, 2009, to win cash prizes up to Rs 55,000!

After the successful completion of the "Training Conference" in Bangalore, SiliconIndia is now coming to **Hyderabad**

It's not worth a miss If you are

Just out of the College, Looking for change in career, thinking about what to do after completion of engineering degree... Many question whether to do specialization on Networking, Mainframe, SAP, Oracle, Microsoft certified, PHP, SQL, Java, Cisco networking, hardware, Sun Solaris, VLSI, Embedded and so on.

SiliconIndia's Training Conference is bringing all specialized training Institutes at one platform where you can talk to them and get guidance from their expertise about your career.

SiliconIndia's Training Conference is also having the conference, where experts from Industry will be speaking on the Technologies which you want to opt for your career. This session will be on different domains and areas where you can directly Interact with these expertise through your questions also.

Come now and register yourself today to choose the right career path.

For more Details

you can call at **080 43402000** or can mail at **career@siliconindia.com**

For having the stall call at **080 43402003** or can mail at **alok@siliconindia.com**

Few participating
institutions:



Hyderabad

December 05, 2009 at Sri Sathya Sai Nigamagmam

Sri Nagar Colony, Panjagutt, Hyderabad

Register now for free @ **www.siliconindia.com/events**

Private Cloud Computing

with UEC

Ubuntu Enterprise Cloud (UEC) is built using a number of open source technologies, which include KVM, QEMU, libvirt and Eucalyptus. This article describes the key concepts behind these technologies, and shows how UEC can be used to build a private cloud.



A cloud is a cluster of servers, where servers can be added and removed, on demand. The ability of the cloud to match computing capacity with demand (referred to as elasticity), allows a cloud to offer high performance and scalability, while maintaining high server utilisation, since hardware resources are more closely related with demand and workloads. A cloud thus eliminates the need to maintain excess server capacity.

Clouds are typically of two kinds—public clouds and private clouds. In a public cloud, the servers are owned and managed by a hosting provider (for example, Amazon, Slicehost, etc), and are located at the hosting provider's site. In a private cloud, the servers

are owned and managed by the company/organisation using it, although they can be located either at the company's site or co-located at a hosting provider's site.

Ubuntu Enterprise Cloud (UEC), available in Ubuntu Server 9.10 (Karmic Koala), offers the functionality for both private and public clouds. UEC allows for the creation of a private cloud consisting of servers running Ubuntu Server 9.10. UEC also allows the easy creation of Ubuntu server images that can be deployed on Amazon EC2 public clouds. UEC is built using a number of open source technologies, which include KVM, QEMU, libvirt and Eucalyptus. This article describes the key concepts of these technologies, and shows how UEC can be used to build a private cloud.

Terminology

In order to better describe the underlying concepts and functionality of UEC, it would be best to define a set of common terminologies. Given below are some of the terms used in this article, along with their definitions:

- **Node:** A node refers to a single server or physical machine.
- **Virtual resource:** This is a proxy or abstraction for the real resource. The virtual resource has the same interface and functions as the real resource, but with different attributes. Virtual resources include virtual CPUs, memory, disks and the network.
- **Server virtualisation:** Server virtualisation turns a node into a set of virtual machines (VM), with each VM appearing to have its own set of virtual resources, such as a virtual CPU, memory, disk and network.
- **Hypervisor or a Virtual Machine Monitor (VMM):** A hypervisor or virtual machine monitor is a layer of software that virtualises a node into a set of virtual machines. The hypervisor is responsible for the management of the virtual machines.
- **Virtual Machine (VM):** This is an isolated working environment with its own set of virtual resources, including the CPU, memory, disk and the network. A VM can run an OS inside it. On a given node, it is possible to run one or more VMs.
- **Guest OS:** The guest OS is the operating system that runs within the virtual machine.
- **Domain:** A domain is an instance of an operating system running on a virtual machine. The terms 'domain' and 'virtual machines' are often used interchangeably.
- **Virtual Machine Image:** A virtual machine image, or simply the image, is a file on the hard disk that can be used to create an instance of the virtual machine.
- **Volume:** A volume is either a block device, a raw file or a special format file.
- **Pool:** This provides a means of taking a chunk of storage and dividing it into volumes. A pool can be used to manage a physical disk, NFS server, iSCSI target, host adapter and LVM group.
- **Full virtualisation:** In a full virtualisation solution, the hypervisor must intercept privileged instructions from the guest OS, and these are then simulated by the hypervisor to fulfil the request on the CPU hardware. Simulating instructions inside the hypervisor causes performance degradation. In a full virtualisation solution, no modifications are needed to the guest OS.
- **Para virtualisation:** In a para virtualisation solution, privileged instructions can be run directly against the CPU hardware. The guest OS must be modified in order to cooperate with the underlying hypervisor. A para virtualisation solution is faster than full virtualisation, since no simulation of instructions by the hypervisor is required.

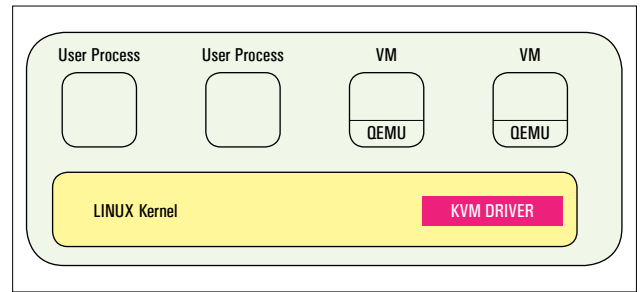


Figure 1: The KVM architecture

- **Hardware-assisted virtualisation:** A hardware-assisted virtualisation solution requires that the underlying CPU hardware provide support for virtualisation. This approach combines the benefits of full virtualisation and para virtualisation. No modification is needed to the guest OS, and privileged instructions are executed by the CPU hardware.

KVM

The two dominant virtualisation technologies in the Linux world are Xen and KVM (Kernel-based Virtual Machine)—both being open source. Xen is a para virtualisation-based hypervisor, while KVM is a hardware-assisted virtualisation-based hypervisor. UEC is based on KVM.

KVM is a hardware-assisted virtualisation solution for Linux on x86 hardware containing virtualisation extensions (Intel VT or AMD-V). KVM adds virtualisation support directly to the Linux kernel. It was introduced to the Linux kernel version 2.6.20—the current kernel version is 2.6.31.

KVM is essentially only a loadable kernel module (*kvm.ko*) that provides the core virtualisation infrastructure. This is besides the processor-specific module—*kvm-intel.ko* or *kvm-amd.ko*, depending on whether it's an Intel or an AMD processor, respectively. KVM also requires a modified QEMU that runs in user space (we'll discuss this later). KVM itself is responsible for virtualising the memory, while QEMU is responsible for virtualising the I/O.



Note: QEMU is an open source processor emulator. All I/O requests (disk and network I/O) made by the guest OS are intercepted and routed to be emulated by QEMU.

KVM requires virtualisation support from the underlying processor hardware. In the case of Intel processors, Intel-VT capable processors provide virtualisation support (shows in */proc/cpuinfo* as the *vmx* flag). In the case of AMD processors, AMD-V capable processors provide virtualisation support (shows in */proc/cpuinfo* as the *svm* flag). See the *Reference* section at the end of this article for a

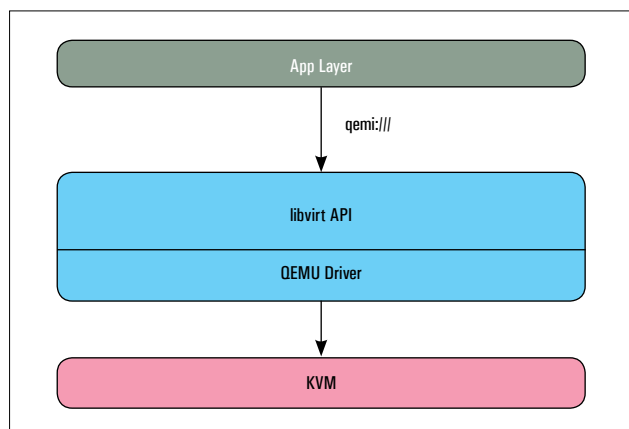


Figure 2: Accessing a local VM using the libvirt API

list of processors supporting virtualisation.

Using KVM, you can run multiple virtual machines, where each VM is a separate Linux process, and is scheduled by the Linux scheduler just like any other process. Each virtual machine can run an unmodified Linux (or Windows) image. Each of these virtual machines have private virtualised hardware—viz., a network card, disk, graphics adapter, etc.

Note that any normal Linux process has two modes: the kernel mode and user mode. In kernel mode, privileged instructions are executed within the kernel. In user mode, non-privileged instructions are executed in user space. KVM adds a third mode called the guest mode, and certain instructions of the guest OS are executed in this mode.

For more details on how to work with KVM, please refer to: www.linuxforu.com/how-to/kvm-virtualisation-the-linux-way. Figure 1 gives an architectural overview of KVM.

JeOS

Now that we're done with an overview on KVM, let's get familiar with another term—Ubuntu JeOS (just-enough operating system).

Quoting the Ubuntu Documentation [help.ubuntu.com/community/JeOS]: “Ubuntu JeOS (pronounced ‘Juice’) is an efficient variant of the Ubuntu Server OS, configured specifically for virtual appliances. It comes with a tuned kernel that only contains the base elements needed to run within a virtualised environment. Ubuntu JeOS excludes unnecessary drivers, and retains only the minimal required packages.”

You can create an Ubuntu JeOS-based VM using a program called `ubuntu-vm-builder`—described in more detail in the next section.

Installing KVM

The first step is to install KVM and associated programs:

```
$ sudo aptitude install kvm libvirt-bin \
ubuntu-vm-builder bridge-utils
```




Figure 3: Accessing a remote VM using the libvirt API

The following are the package descriptions:

- *kvm* is the KVM driver.
- *libvirt-bin* provides *libvirtd* which is needed to administer QEMU and KVM instances using *libvirt*.
- *ubuntu-vm-builder* is a powerful command-line tool for building JeOS virtual machines.
- *bridge-utils* provides utilities to create a bridge from the virtual machines to the physical network interface (e.g., an Ethernet card).

A default virtual bridge interface *virbr0* is created automatically when installing KVM. In case bridged-networking is required, a *br0* bridge will need to be created in `/etc/network/interfaces`.

The next step is to create a JeOS virtual machine using *ubuntu-vm-builder* as follows:

```
$ sudo ubuntu-vm-builder kvm karmic --arch 'amd64' \
--mem '1024' --rootsize '10240' --swapspace '2048' \
--kernel-flavour 'server' --hostname 'vb-euc1' \
--mirror 'http://archive.ubuntu.com/ubuntu' \
--components 'main,universe' --name 'admin' \
--user 'admin-server' --pass 'welcome123' \
--ip '192.168.0.20' --libvirt qemu:///system ;
```

Here:

- *vm* is the target virtualisation software—in this case, KVM.
- *suite* is the specified Ubuntu suite (Karmic).
- *arch* is the target architecture (amd64, i386).
- *mem* is the memory in MB, of the RAM (1024 MB).
- *rootsize* is the disk space for the root filesystem (10240 MB = 10GB).
- *swapspace* is the swap space (2 GB).
- *kernel-flavour* is the kernel image flavour to install (Ubuntu server).

SharePoint Comes Back to San Francisco!

Attend



SPTechCon

The SharePoint
Technology Conference

Feb. 10-12, 2010 → San Francisco

"Wide range of topics—good depth—plus it was fun!"

Patrick McMahon, Web/IS Architect,
SEPATON Inc.

"Excellent content and speakers.
Knowledgeable, real-world
examples gave me solutions
to various SharePoint problems."

Laura Diorio, Web Systems Manager, Porter Novelli

Choose From
**80+ Classes
& Workshops**

SPECIAL KEYNOTES



Ted Pattison
Co-founder,
Critical Path
Training




Tom Rizzo
Director of
SharePoint,
Microsoft


DIAMOND SPONSOR




PLATINUM SPONSORS



GOLD SPONSORS



SILVER SPONSORS



Check Out Our All-Star Faculty!



Shane
Young



Randy
Drisgill



Todd
Klindt



John
Ross



Michael
Noel



Bob
Mixon



Heather
Solomon



Andrew
Connell



Bill
English



Dux
Raymond Sy



Joel
Oleson



Errin
O'Connor



Laura
Rogers



Mauro
Cardarelli



Mark
Miller



Matt
Passannante



Mark
Rackley



Jason
Dearing



Eric
Harlan



Shadeed
Eleazer



Phillip
Wicklund



Dan
Usher



Joshua
Haebets



Peter
Serzo



Nicola
Young



Mark
Ferraz



Karuana
Gatimu



Steven
Fowler



Paul J.
Swider



Jennifer
Mason

For more information, go to
www.sptechcon.com



REGISTER NOW

for Early Bird Rates!
SAVE \$\$\$!

A BZ Media Event

February 10-12, 2010

Hyatt Regency

San Francisco Airport



- *ip* is the IP address, in dotted form, of the virtual machine to be created.
- *hostname* and *domain* configures the hostname and domain name of the VM.
- *user*, *name*, and *pass* set the user name, full name, and password of the default user.
- *libvirt qemu:///system* is added to ensure that we will use *virsh* to manage the VM.

Virtual machine management

libvirt is a generic, virtualisation technology-independent API to securely manage local as well as remote domains or VMs. It allows one to create, start, stop, modify, monitor and migrate VMs.

The *libvirt* API provides support for different types of hypervisors. In order to access a VM, either local or remote, a URI is specified, along with the type of hypervisor (*libvirt*'s driver and URI approach). This ensures the correct VM is accessed.

libvirtd is a daemon to manage QEMU VM instances and *libvirt* virtual networks. *libvirtd* is configured by */etc/libvirt/libvirtd.conf*.

virsh is a command-line utility that you can use to issue *libvirt* API calls. Some of the commonly used *virsh* commands are:

- *list* – to list running virtual machines:
`virsh -c qemu:///system list`
- *start* – to start a virtual machine
`virsh -c qemu:///system start kvm-euc-demo-server-karmic`
- *autostart* – to start a virtual machine at boot (makes it an upstart job):
`virsh -c qemu:///system autostart kvm-euc-demo-server-karmic`
- *reboot* – to reboot a virtual machine:
`virsh -c qemu:///system reboot kvm-euc-demo-server-karmic`
- *save* – to save the state of a virtual machine to a file in order to be restored later (once saved, the VM will no longer be running):
`virsh -c qemu:///system save kvm-euc-demo-server-karmic \`
`kvm-euc-demo-server-karmic.state`
- *restore* – You can restore a saved virtual machine using:
`virsh -c qemu:///system restore kvm-euc-demo-server-karmic.state`
- *shutdown* – to shutdown a virtual machine:
`virsh -c qemu:///system shutdown kvm-euc-demo-server-karmic`

libvirt stores its configuration as XML in */etc/libvirt/qemu*. While it is possible to modify the XML directly, it is better to modify it using *virsh* or *virt-manager*. It is possible to add CPU, memory, add/change a network card used by a VM, etc.

Virtual networking

There are essentially two different ways for a VM to access the network. The first is to use NAT forwarding or virtual networks. This is provided by default, and traffic is 'NATed' through a virtual bridge device *virbr0* to the

Eucalyptus and the various controllers

In order to understand Eucalyptus better, we would recommend you read an excellent paper called *The Eucalyptus Open-source Cloud-computing System*, available at <http://open.eucalyptus.com/wiki/Presentations>. This paper explains the various controllers as follows:

- Node Controller (NC) controls the execution, inspection, and termination of VM instances on the host where it runs.
- Cluster Controller (CC) gathers information about and schedules VM execution on specific node controllers, and manages the virtual instance network.
- Storage Controller (Walrus) (SC) is a put/get storage service that implements Amazon's S3 interface, providing a mechanism for storing and accessing virtual machine images and user data.
- Cloud Controller (CLC) is the entry-point into the cloud for users and administrators. It queries node managers for information about resources, makes high-level scheduling decisions, and implements them by making requests to cluster controllers.

network. In this case the VM does not connect directly to the LAN; rather, it connects through the virtual bridge using NAT forwarding. The commands *virsh net-list -all* and *brctl show* can be used to see the details of *virbr0*.

The second way for a VM to access the network is using bridge networking or a shared physical device. The bridge *br0* is configured in */etc/network/interfaces*, and this allows the VM to connect directly to the LAN.

Virtual Machine Manager

The Virtual Machine Manager is a graphical utility to manage local and remote virtual machines. It is part of the *virt-manager* package (installed using *sudo apt-get install virt-manager*).

To connect to the local *libvirt* service, use the following command:

```
virt-manager -c qemu:///system
```

To connect to the *libvirt* service running on a remote host, use:

```
virt-manager -c qemu+ssh://virtnode1.mydomain.com/system
```

UEC and Eucalyptus

The private cloud functionality of UEC is provided by Eucalyptus (short for Elastic Computing Architecture for Linking Your Programs to Useful Systems), which is an open source system that helps to transform a group of Linux servers into a cloud computing platform.

MODEL ENGINEERING COLLEGE IN ASSOCIATION WITH GULF OIL

PRESENTS

EXCEL X

INSPIRE | INNOVATE | ENGINEER

TEDx MEC
x = independently organized TED event

IN THE SPIRIT OF IDEAS WORTH SPREADING, TED HAS CREATED A PROGRAM CALLED TEDx. TEDx IS A PROGRAM OF LOCAL, SELF-ORGANIZED EVENTS THAT BRING PEOPLE TOGETHER TO SHARE A TED-LIKE EXPERIENCE. OUR EVENT IS CALL TEDx (PLACE), WHERE X-INDEPENDENTLY ORGANIZED TED WILL COMBINE TO SPARK DEEP DISCUSSION AND CONNECTION IN A SMALL GROUP. THE TED CONFERENCE PROVIDES GENERAL GUIDANCE FOR THE TEDx PROGRAM, BUT INDIVIDUAL TEDx EVENTS, INCLUDING OURS, ARE SELF-ORGANIZED.

WE INVITE YOU TO BE A PART OF THIS EVENT HOSTED FOR THE FIRST TIME BY AN INDIAN ENGINEERING COLLEGE. SEE YOU THERE ON DECEMBER 13TH! VISIT [HTTP://TEDX.MEC.AC.IN/](http://tedx.mec.ac.in/) FOR MORE INFO

DECEMBER
12 13 13

HACK FEST

PRINCIPAL MATHEMATICS

WHAT'S THE FUNDO?

MANAGERIAL METTLE

LOAD OF THE CODE

PAPYRUS OF AN

WAVE CLONING

EXTANSICITY

ROBO COMBAT

TELE CODER

RO-BOLT

DEAFUSE

TAP

SLASH BIN SLASH BASH

TEDx MEC

PRIZES WORTH
RS 3 LAKH

EVER WONDERED WHO CODED YOUR FAVOURITE VIDEO PLAYER? EVER THOUGHT HOW YOU GET THE COMPLETE SUITE OF SOFTWARES FREE OF COST? EVER STUMBLED ON BUGS IN YOUR FAVOURITE SOFTWARE AND SAID WTF!! NOW, STOP COMPLAINING, START CODING. HACKFEST IS ONE SMALL DROP OF IT. THE EVENT AIMS TO SERVE AS AN INTRO-DUCTORY EVENT FOR CONTRIBUTING TO THE FREE WORLD.

WE TAKE SOFTWARE THAT WE USE IN OUR DAY TO DAY LIFE AND FIX BUGS, ADD FEATURES AND MORE. THE EVENT WILL FOCUS MAINLY ON USABILITY ENHANCEMENTS AND BUGS SO AS TO PROVIDE EVERYONE A QUICK STEP INTO THE OPEN CULTURE AND SHARING.

WE TAKE SOFTWARE THAT WE USE IN OUR DAY TO DAY LIFE AND FIX BUGS, ADD FEATURES AND MORE. THE EVENT WILL FOCUS MAINLY ON USABILITY ENHANCEMENTS AND BUGS SO AS TO PROVIDE EVERYONE A QUICK STEP INTO THE OPEN CULTURE AND SHARING.

www.excelmec.org



OEN CONNECTORS



SINCERELY YOURS

Eucalyptus was developed at the Computer Science Department of the University of California at Santa Barbara as a research project, and is now enhanced and maintained by Eucalyptus Systems, a company founded by the original developers.

A UEC cloud consists of one or more nodes, where each node has one node controller. One or more node controllers are managed by a cluster controller, and the entire cloud is managed by a cloud controller. Refer to the box titled 'Eucalyptus and the various controllers' to get an idea about the architecture.

Installing Eucalyptus

This article describes the installation of UEC on a two-machine configuration, in which Node 1 contains the CLC, CC and SC, while Node 2 contains the NC. Other variations include having the CLC, CC, and SC on separate nodes, as well as having multiple NC and multiple nodes.

First, we install the CLC, CC and SC packages on Node 1:

```
$ sudo apt-get install eucalyptus-cloud eucalyptus-cc \
eucalyptus-walrus eucalyptus-sc
```

Next, we install the NC package on Node 2:

```
sudo apt-get install eucalyptus-nc
```

We configure `/etc/eucalyptus/eucalyptus.conf` with the name of the bridge (`virbr0`), and restart the NC.

We install the CLC Eucalyptus user's public SSH key into the NC Eucalyptus user's `authorized_keys` file.

We need to register the CC and SC with the CLC. Since, in this installation, the CC and SC are on the same node as the CLC, the registration is automatic. In case they are on different nodes, the registration is accomplished using the following code:

```
$ sudo start eucalyptus-cc-registration
$ sudo start eucalyptus-sc-registration
$ sudo start eucalyptus-walrus-registration
```

The next step is to register the NCs with the CLC. This is done by adding the IP address of each NC in `/etc/eucalyptus/eucalyptus.conf` of the CLC.

We next need to obtain credentials on the CLC. This is done using the code below:

```
$ sudo euca_conf --get-credentials mycreds.zip
```

The credentials (`unzip mycreds.zip`) are stored in the `.euca` directory in the home of the Eucalyptus user.

Euca2ools are installed using the code below:

```
$ sudo apt-get install euca2ools
```

To check if things are working fine, issue the following command:


```
$ euca-describe-availability-zones verbose
```

This should list the details of the cloud.

Finally, we need to install an Ubuntu Karmic image to run in the virtual machines. This is done by connecting to the following:

<https://<cloud-controller-ip-address>:8443/>

Log in for the first time using `admin/admin`, and then change the password. After image creation is complete, the image is ready to run in the VM. Please refer to the UEC package install link in the References section below for more details on the steps to install UEC.

That's all folks—you now have your private cloud ready for use.  **END**

References

- KVM: http://www.linux-kvm.org/page/Main_Page
- Ubuntu KVM install: <https://help.ubuntu.com/community/KVM>
- Ubuntu KVM install Howto: <http://www.howtoforge.com/virtualization-with-kvm-on-ubuntu-9.04>
- QEMU: <http://www.qemu.org/>
- libvirt: <http://libvirt.org/>
- Eucalyptus: <http://open.eucalyptus.com/>
- UEC: <https://help.ubuntu.com/community/UEC>
- UEC package install:
- Processors supporting virtualization: http://en.wikipedia.org/wiki/X86_virtualization
- UEC package install: <https://help.ubuntu.com/community/UEC/PackageInstall>

By: Ashwin Mangale and Amit Mathur

The authors are part of the leadership team at Vector Brook (www.vectorbrook.com), a Pune-based software company, where they are involved in the development of cloud-based applications and solutions for the education and enterprise markets.

Asia's Leading Open Source Magazine

To BUY or SUBSCRIBE, please visit www.linuxforu.com

www.linuxforu.com





EveryWhere & AnyWhere

32 lakh active young sports loving men read
cricket today every month*

Lifeline of Cricket

Cricket
today

www.crickettoday.com

*SOURCE IRS 2008

Available in English & Hindi



Tips & Tricks



Message of the day

Normally when you log into a console you aren't greeted by any message (except in openSUSE, which always tells you to "Have a lot of fun!"). If you want to be greeted by a message the moment you log in, you can add that message in the `/etc/motd` file.

—Vaseem Javeed, vaseemjaveed@gmail.com



The CPU isn't busy; it's time to run

If you want to run a program or command when the CPU is free or not very busy, the tool you need is `batch`. It only executes commands when the system load levels permit. In other words, when the load average drops below 1.5, or the value specified in the invocation of `atd`.

```
echo 'tar czf /tmp/mydocs.tar.gz /home/me/docs/' | batch
```

...or enter:

```
batch
```

...and the following prompt appears at your console:

```
at>
```

You can type the command and press Ctrl + d to save the job.

—Remin Raphael, remin13@gmail.com



Mounting NTFS partition on RHEL

You should first install a package called `ntfs-3g`. Open the terminal and type:

```
yum install ntfs-3g
```

To auto mount the NTFS partition, add a line to

```
/etc/fstab:
```

```
/dev/sda2 /mnt/ ntfs defaults 0 0
```

...followed by:

```
mount -a
```

Your NTFS partition should now become accessible.

—Vijith P A, vijith.pa@gmail.com



Selecting alternative apps in Debian

In order to change the default Web browser of Debian GNU/Linux, you can use the `update-alternatives` command as follows:

```
$ sudo update-alternatives --config x-www-browser
```

There are 4 alternatives which provide 'x-www-browser'.

Selection	Alternative
-----------	-------------

1	/usr/bin/epiphany
2	/usr/bin/iceweasel
*+	3 /usr/bin/konqueror
4	/usr/bin/epiphany-gecko

Press enter to keep the default[*], or type selection number:

Note that the asterisk sign indicates Konqueror is currently the default browser. In order to switch the

default to Iceweasel I'll press '2'. Here's how the console tells us that the change has been configured:

Using '/usr/bin/iceweasel' to provide 'x-www-browser'.

Note that the above command requires root privileges.

Similarly, if you want to choose your default Java installation:

```
$ sudo update-alternatives --config java
```

There are 3 alternatives which provide 'java'.

Selection	Alternative
1	/usr/bin/gij-wrapper-4.1
*+ 2	/usr/lib/jvm/java-gcj/jre/bin/java
3	/usr/lib/jvm/java-1.5.0-sun/jre/bin/java

Press enter to keep the default[*], or type selection number: 3

Using '/usr/lib/jvm/java-1.5.0-sun/jre/bin/java' to provide 'java'.

Note that we switched from java-gcj to Sun's Java. Let's check if it has indeed become the default:

```
$ java -version
java version "1.5.0_14"
Java(TM) 2 Runtime Environment, Standard Edition (build 1.5.0_14-b03)
Java HotSpot(TM) Client VM (build 1.5.0_14-b03, mixed mode, sharing)
```

—Senthil Kumaran S, stylesen@gmail.com



Retrieving data from the RAM

You're working on an important document and the following happens:

- The application crashes!
- You close it without saving
- You save over a document you had just opened!

Solution:

While you've lost what you had on screen, you may have forgotten that the computer still has it stored in the RAM. You just need to find a way to get it.

First things first—every step you take after the realisation of the loss can cause pieces of it to get overwritten (since this memory is now marked as free by the OS). So the first step is to create a snapshot of your current RAM. If you have a lot of RAM, this can

take over 1 GB of space on your hard drive.

```
cp /dev/mem ~/memory.bin
```

After it's done copying, you can access the file using *less*:

```
less ~/memory.bin
```

When the program starts, it'll report that it's trying to compute line numbers. Skip this by hitting *Ctrl+c* once.

Now search for a unique word that you remember, typing in your document by hitting the forward slash (/) followed by the word (e.g., /trickle). After much churning, matches will be displayed. Hit the up and down arrow keys to scroll from the hit point. To continue searching for the same phrase, simply hit '/' and Enter again. Once you stumble upon your document, it may be in bad shape, but there should be enough information left to be copied out of the terminal.

Note:

1. You must be logged in as the root to read the mem device.
2. You will save a lot of file space and make your job easier if you use this line, *cat /dev/mem | strings > ~/memory.bin*

—Bharat kumar, bharatrayudu@gmail.com



Creating secure passwords

Time and again we need to create strong passwords for Web forms, e-mail accounts, Web registration pages, etc. We can create one easily in GNU/Linux using the *openssl* command as follows:

```
openssl rand 12 -base64
```

The above command will create a random base 64 encoding string each time it's run. Since the string you get from the command is different each time the command is run, it's secure and easy to create a strong password this way.

—Swapneel, swapneel@eth1.in



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.



Performance Tuning and Monitoring **Part 3**

Here're some tricks to help you counter disk fragmentation and improve RAID performance on hard drives significantly.

In the October issue of LFY, we discussed the different performance tuning methods related to hard disk drives, using IO elevators. We had also explored the queuing theory and finding hotspots using the strace command. Now, after our Diwali break, it's time we resume fine tuning our hard disks again.

This time we will look at how to counter disk fragmentation and improve RAID performance significantly. We will also learn about cache profiling using Valgrind.

Countering disk fragmentation

Disk fragmentation affects the sequential read access performance. This is because it results in extra head movement on the disk drive. In fact, the Linux filesystem is designed to minimise the effects of fragmentation as much as possible. When a file is created or extended, the filesystem always tries to allocate blocks for the file from the same block group that contains its inode. And to ensure that files are spread evenly in the

partition, when a new file is created, an attempt is made to place it in some block group other than the one containing its parent directory.

To view the fragmentation of a file we can use the following command:

```
$ filefrag -v /path/to/the/filename
```

To check the fragmentation on a mounted filesystem, issue the command shown below:

```
$ dumpe2fs /dev/sdaX
```

...where X is the partition number.

Types of journaling in ext3

One major issue in journaling filesystems is whether they only log changes to the filesystem metadata or log changes to all filesystem data, including changes to the files themselves. The ext3 filesystem supports three different journaling modes. Here's how Wikipedia [<http://en.wikipedia>.

[org/wiki/Ext3](#)] explains each:

- **Journal (lowest risk):** Both metadata and file contents are written to the journal before being committed to the main filesystem. Because the journal is relatively continuous on disk, this can improve performance in some circumstances. In other cases, performance gets worse because the data must be written twice—once to the journal, and once to the main part of the filesystem.
- **Ordered (medium risk):** Only metadata is journaled; file contents are not, but it's guaranteed that file contents are written to disk before associated metadata is marked as committed in the journal. This is the default on many Linux distributions.
- **Writeback (highest risk):** Only metadata is journaled; file contents are not. The contents might be written before or after the journal is updated. As a result, files modified right before a crash can become corrupted.

These modes can be made permanent by editing the `/etc/fstab` file and specifying the journaling mode, as shown in Figure 1.

Improving journal performance

By default, the journal is located at an inode within the filesystem. To improve the performance, we should try to keep the journal on a separate device. The journaling filesystem must use the same block size as that of the data filesystem. You should always dedicate the entire partition to the journal and make sure that the external journal is always in a disk drive capable of equal or better performance.

I'm creating a normal RAID-5 device here, and then I will tune the array by using the STRIDE and CHUNK size. Finally, I'll move the journal to an external device to enhance performance.

Creating a normal RAID-5 array

Step 1 Create three partitions of 100 MB (that is, 102400 KB) each

sysfs	/sys	sysfs	defaults	0 0
proc	/proc	proc	defaults	0 0
/dev/sda5	/data	ext3	defaults,data=ordered	0 0
/dev/sda6	/vbox	ext3	defaults	0 0

Figure 1: Changing the journaling mode of an ext3 filesystem by editing `/etc/fstab`

[root@server1 ~]# iostat -x -d /dev/sda 1												
Linux 2.6.18-92.el5xen (server1.example.com) 09/14/2009												
Device:	rrqm/s	wrqm/s	r/s	w/s	rsec/s	wsec/s	avgrq-sz	avgqu-sz	await	svctm	%util	
sda	9.81	11.40	10.40	4.12	1461.35	123.99	109.20	0.06	4.24	2.20	3.20	
Device:	rrqm/s	wrqm/s	r/s	w/s	rsec/s	wsec/s	avgrq-sz	avgqu-sz	await	svctm	%util	
sda	49.00	0.00	39.00	0.00	7816.00	0.00	200.41	0.10	2.56	2.26	8.80	

Figure 2: Calculating the chunk size

and change its 'type' to `fd`.

Step 2 Use the `mdadm` command to create the array:

```
# mdadm -C /dev/md0 -l 5 -n 3 /dev/sda[7,8,9]
```

Step 3 Format the array:

```
# mke2fs -b 2048 -j /dev/md0
```

Step 4 The RAID device is now ready and initiated. You can mount it on any location to use it.

Step 5 I will now capture the filesystem layout using the `dumpe2fs` command:

```
# dumpe2fs /dev/md0 > anyfilename
```

Step 6 Now stop the array and its superblock:

```
# mdadm -S /dev/zero --zero-superblock
```

I now want to reduce the disk visit count on this array by using the chunk size and stride values.

Chunk size is actually the amount of data read/written from each device in an array before moving to the next device in a round-robin manner. It is also known as 'granularity of the stripe'. It's always in KB. Chunk size is very important for RAID performance.

If the chunk size is too small, it can result in a file spreading into multiple smaller pieces thus increasing the number of drives it will use. On the other hand, if the chunk size is too big, a situation may develop in which all the IO is handled by one physical device only—a condition

Inode count:	25624576
Block count:	25599577
Reserved block count:	1279978
Free blocks:	18711564
Free inodes:	25582961
First block:	0
Block size:	4096
Fragment size:	4096
Reserved GDT blocks:	1037

Figure 3: A snippet of `dumpe2fs` output

known as 'hot spot'.

The ideal situation is to spread the IO across all the devices evenly.

To calculate the chunk size, first use the following code:

```
# iostat -x -d /dev/sda 1
```

...to get the value of `avgrq-sz` (average request size). Figure 2 displays the output on my system.

I will consider the value of 200.41 for my calculations. Just note down this value. Now divide it by 2. It's approximately 100.20.

Now divide 100.20 by the number of disk drives in your array. In my case, it was three. So I should be using the figure 100.20/3, which equals 33.4.

Round off this value to the nearest power of 2, which comes to 32. That means my chunk size is 32.

Now, before we move ahead in configuring a tuned RAID, we should also find out the stride value. Stride is the same as chunk size but uses a different value. The aim of the stride is to distribute the block bitmap across RAID member devices.

The formula to calculate the stride is: $\text{STRIDE} = \text{CHUNK SIZE} / \text{BLOCK SIZE}$

We have already retrieved the block size from the `dumpe2fs /dev/md0` command by redirecting the output to the `anyfilename` file. Now, access this file with the `less`

```
[root@legacy alok]# x86info -c
x86info v1.24. Dave Jones 2001-2009
Feedback to <davej@redhat.com>.

Found 2 CPUs
-----
CPU #1
EFamily: 0 EModel: 0 Family: 6 Model: 15 Stepping: 6
CPU Model: Core 2 Duo [B2]
Processor name string: Intel(R) Core(TM)2 CPU          T5500   @ 1.66GHz
Type: 0 (Original OEM) Brand: 0 (Unsupported)
Number of cores per physical package=2
Number of logical processors per socket=2
Number of logical processors per core=1
APIC ID: 0x0 Package: 0 Core: 0 SMT ID 0
Cache info
L1 Instruction cache: 32KB, 8-way associative. 64 byte line size.
L1 Data cache: 32KB, 8-way associative. 64 byte line size.
L2 cache: 2MB, 8-way associative. 64 byte line size.
TLB info
Instruction TLB: 4x 4MB page entries, or 8x 2MB pages entries, 4-way associative
Instruction TLB: 4K pages, 4-way associative, 128 entries.
Data TLB: 4MB pages, 4-way associative, 32 entries
L1 Data TLB: 4KB pages, 4-way set associative, 16 entries
L1 Data TLB: 4MB pages, 4-way set associative, 16 entries
Data TLB: 4K pages, 4-way associative, 256 entries.
64 byte prefetching.
-----
CPU #2
EFamily: 0 EModel: 0 Family: 6 Model: 15 Stepping: 6
CPU Model: Core 2 Duo [B2]
Processor name string: Intel(R) Core(TM)2 CPU          T5500   @ 1.66GHz
Type: 0 (Original OEM) Brand: 0 (Unsupported)
Number of cores per physical package=2
Number of logical processors per socket=2
Number of logical processors per core=1
```

Figure 4: Cache information using the x86info tool

command—it's a huge one—and look for the part shown in Figure 3.

So my stride would be 32/4, or 8.

Now, let's create the tuned array using the chunk size and stride value. We'll use the `mdadm` command as shown:

```
# mdadm -C /dev/md0 -l 5 -n 3 -C 32 /dev/sda7
/dev/sda8 /dev/sda9
```

Now format the device using the stride value. For an ideal filesystem, you can assume a block size of 2 KB (that is, 2048 bytes). But here I am using the actual block size of my filesystem, which is 4 KB or 4096 bytes.

```
# mke2fs -b 4096 -E stride=8 -j /dev/md0
```

That's it! Our job is done. Now just mount the tuned RAID array to some directory (`/raid-location`). You can use the `/etc/fstab` file to make it permanent.

Mounting the journal to the external filesystem

Till now the journal has been on the same filesystem that can lead

to performance degrade. Naturally, it's always better to mount the FS journal to some external or separate filesystem.

Step 1 First *umount* the RAID device:

```
# umount /raid-location
```

Step 2 Check the block size of the RAID filesystem using the `dumpe2fs /dev/md0` command.

Step 3 Create a partition using the `fdisk` command that has to be used for the journal in future.

Step 4 Format the new partition that you had just created to use as the journal device:

```
# mke2fs -O journal_dev -b 4096 -L raid_journal
/dev/sda10
```

Step 5 Before using the new journal filesystem, remove the existing journal from `/dev/md0`:

```
# tune2fs -O ^has_journal /dev/md0
```

Step 6 Now since the old journal has been removed from `/dev/md0`, the new partition `/dev/sda10` can be used to hold the journal for `/dev/md0`:

```
# tune2fs -J device=/dev/sda10 /dev/md0
```

Step 7 Mount the RAID device now:

```
# mount /raid-location
```

Locality of reference

CPU cache (L1/L2/L3 cache) is a very important factor in performance. Its speed is normally twice the speed of RAM and is used to hasten the memory retrieval process. Without a cache the life of the CPU will be like that of librarians without a storage-shelf behind them. So every time people request a book or need to store it, they have to climb up the ladder to place the book or get the book from the main shelf (main memory) itself. That will obviously slow down the librarian's performance and lower the customer's satisfaction levels.

The cache is divided into a few lines. Each line is used to cache a specific location in the memory. We have separate cache instructions for the processor: I-cache, data instructions, and D-cache. This classification is also known as the Harvard Memory Architecture.

Every cache has an associated cache-controller. So whenever the CPU asks for any reference from the main memory, the cache-controller first checks if that requested reference is in the cache. If the reference is found there, the request is catered to from there itself, without going to the main memory. So this significantly increases performance. This is known as cache-hit.

However, if the request is not found there (known as cache-miss), its requested location is brought from the main memory to the cache for reference (known as cache-line-fill). On a multi-processor system, if one CPU cache gets a cache-line-fill, then the first CPU must inform the second CPU about this. So the other CPU

India's only Computer Networking Institute by Corporate Trainers.

Network NUTS

for networking, for life

Happy New Year 2010

Rs. 2010/- Discount on all Courses.

**"RH436
(Red Hat Enterprise
Clustering and Storage
Management) batch
starting from
21st-24th Dec.09"**

RHCA*

RHCSS

RHCE

Open LDAP

Virtualization

Shell Scripting

DansGuardian Administration

OpenDNS Administration

Nagios Server Administration

Fact #1

Given 1830+ Red Hat Certified professional (RHCE + RHCSS), with BEST result in Asia-Pacific region.



Boot Camps - RHS333 : 07th Dec. - 10th Dec.09

RHS423 : 14th Dec. - 17th Dec.09

RHS429 : 21st Dec. - 24th Dec.09

EX333 2nd & 27th Dec.09

EX423 24th Dec.09

EX429 24th Dec.09

RHCE 10th & 29th Dec.09

Thank You all for making us your preferred training partner :

Wartsila India Ltd.

Life Insurance Corporation of India Limited, Govt. of India Enterprise

National Information Centre (NIC), Govt. of India Enterprise

Indian Air Force (WAC)

Indian Air Force (CAC)

Indian Air Force (AFRO)

Indian Air Force (AFCAO)

Bajaj Allianz

Indian Meteorological Department (IMD), Govt. of India Enterprise

Meteo France International, France

DoIT, Govt. of Rajasthan

St. Patricks International College, London, UK

GeoEnpro Petroleum Ltd., India

IFFCO, Govt. of India Enterprise

CRIS (Indian Railways), Govt. of India

IBM

Delhi University Computer Centre, Delhi University

Institute of Himalayan Bio-Resource Technology (Council of CSIR)

Broadcast Engineering Consultant India Limited (Govt. of India Enterprise)

Department of General Hydrocarbons (Govt. of India)

redhat
PREMIER
PARTNER



South Delhi : A-184, Bhisham Pitamah Marg, Opp. Defence Colony, Delhi-2 Tel. : 46526981-2, 9310024503

East Delhi : 508, DDA Complex, Laxmi Nagar Distt. Centre, Delhi-92 Tel. : 42448951-2, 9310024502

North Delhi : 2513, Hudson Lane, Kingsway, Camp, North Campus, Delhi-9 Tel. : 65156734, 9310024501

www.networknuts.net

need not retrieve it from the main memory. This process is known as cache-snooping.

With the theory out of our way, how can we check the cache? Here's a simple command that does that job:

```
# getconf -a | grep -i cache
```

The following is the output from my test system:

```
LEVEL1_ICACHE_SIZE      32768
LEVEL1_ICACHE_ASSOC      8
LEVEL1_ICACHE_LINESIZE   64
LEVEL1_DCACHE_SIZE      32768
LEVEL1_DCACHE_ASSOC      8
LEVEL1_DCACHE_LINESIZE   64
LEVEL2_CACHE_SIZE      2097152
LEVEL2_CACHE_ASSOC      8
LEVEL2_CACHE_LINESIZE   64
LEVEL3_CACHE_SIZE        0
LEVEL3_CACHE_ASSOC      0
LEVEL3_CACHE_LINESIZE   0
LEVEL4_CACHE_SIZE        0
LEVEL4_CACHE_ASSOC      0
LEVEL4_CACHE_LINESIZE   0
```

You can also use the *x86info* tool (you need to install it first). It gives very detailed information about your CPU and types of cache. Figure 4 shows the output on my system when I execute the *x86info -c* command.

You can also check */var/log/dmesg* for this information.

Locality of reference

Applications normally tend to behave in a particular way when accessing data. An application accessing memory location X is more likely to access memory location X+1 in the next few cycles of execution. This behaviour is known as spatial locality of reference.

Programs that access the memory sequentially will generally get more benefits from the cache. Programs that result in more cache-misses are more expensive, as they increase latency. There is a tool called Valgrind that you can use to profile cache usage. Given a choice, this can help you determine the more efficient program.

```
=6941== Cachegrind, a cache and branch-prediction profiler.
=6941== Copyright (C) 2002-2008, and GNU GPL'd, by Nicholas Nethercote et al.
=6941== Using LibVEX rev 1884, a library for dynamic binary translation.
=6941== Copyright (C) 2004-2008, and GNU GPL'd, by OpenWorks LLP.
=6941== Using valgrind-3.4.1, a dynamic binary instrumentation framework.
=6941== Copyright (C) 2000-2008, and GNU GPL'd, by Julian Seward et al.
=6941== For more details, rerun with: -v
=6941==
Starting
finished
=6941==
=6941== I  refs:      6,188,126,666
=6941== I1 misses:      666
=6941== L2i misses:      664
=6941== I1 miss rate:    0.00%
=6941== L2i miss rate:   0.00%
=6941==
=6941== D  refs:      3,937,849,712 (3,375,259,224 rd + 562,590,488 wr)
=6941== D1 misses:      35,157,103 ( 663 rd + 35,156,440 wr)
=6941== L2d misses:      35,157,056 ( 621 rd + 35,156,435 wr)
=6941== D1 miss rate:    0.8%      0.0%      +      6.2% )
=6941== L2d miss rate:   0.8%      0.0%      +      6.2% )
=6941==
=6941== L2 refs:      35,157,769 ( 1,329 rd + 35,156,440 wr)
=6941== L2 misses:      35,157,720 ( 1,285 rd + 35,156,435 wr)
=6941== L2 miss rate:    0.3%      0.0%      +      6.2% )
root@legacy ~]#
```

Figure 5: Testing application performance using Valgrind

Here are the steps on how to get started with it:

Step 1 Use the *x86info -c* command to get the values of your cache. Valgrind is more interested in: Instruction Cache (I1), Data Level 1 (D1) and Data Level 2 (L2) cache values.

Remember that Valgrind demands the I1 cache value in bytes rounded off to the nearest power of 2, besides the D1 and L2 cache (also in bytes). We also need to jot down the values of associativity and line-size. From my output of *x86info -c*, as shown in Figure 4, I get the following values.

- I1 = 32768,8,64
- D1 = 32768,8,64
- L2 = 2097152,8,64

Step 2 Now, if I have two applications that perform the same task, I need to check which one performs better with respect to cache-hits and cache-misses. I can use Valgrind to find out:


```
# valgrind --tool=cachegrind --I1=32768,8,64 \
-- D1=32768,8,64 \
-- L2=2097152,8,64 sample1
```

Here I'm testing a sample

application called sample1. Figure 5 shows the output (truncated) of the above command.

As you can see, the cache-miss rate for the application sample1 is very low (below 1 per cent). This means that sample1 makes good use of cache-hits, thereby enhancing the performance. It is always recommended to use Valgrind for new applications before putting it in a commercial environment.

As a general rule, the less the miss-rate, the better!

Here I conclude for this month. Next month will be dedicated to how we can tune memory address and allocation using PAM, TLB, pdf flush, etc, as well as tune our network performance. **END** 

By: Alok Srivastava

The author is the founder of Network NUTS and holds MCP, MCSE, MCDBA, MCT, CCNA, CCNP, RHCE and RHCSS certifications. Under his leadership, Network NUTS has been a winner of the "Best Red Hat Training Partner in North India" for the last three years in a row. He has also been a trainer for the Indian Air Force, NIC, LIC, IFFCO, Wartsila India Ltd, the Government of Rajasthan, Bajaj Allianz, etc. You can reach him at *alok at networknuts dot net*.



It Plays Music At Your Command (Line)

With CMus, you don't need a memory-hungry GUI music application just to enjoy your music and manage your collection.

After trying out a multitude of GUI music players like Rhythmbox, Exaile, Amarok, and even Quoad Libet, I couldn't find even one that was easy on my old and withered system. Every one of these graphical players was too slow to respond and would cause numerous system freezes. Finally, I decided to give console music players a try.

Initially, I was not willing to go for the console alternatives for many reasons -- the main one being the lack of features and clumsy interfaces of most of these players. Besides, many people are reluctant to use these players due to the steep learning curve involved.

All graphical players have some major strengths, mostly related to the library functions and features like a tree view. Instead of the usual multi-window approach taken by most console music players, the tree browsing saves precious space. Some players come with the capability to display

lyrics and artist information on-the-fly.

So, the main thing you need to prepare yourself for is that a CLI is not a GUI. You will not see album art flying around or any flashy visualisations (though, with some hacking, you can achieve this).

Having gone through almost all the other alternatives, I was not concerned with these shortcomings of CLI players, in general. After some Googling, I decided to give CMus a try--mainly due to its very comprehensive feature set.

The CMus (also called C*) Music Player

CMus music player is what you can call the *vi* of music. It is an ncurses-based music player that uses keybindings of *vi(m)* for navigation and other purposes. Plus, the keys are configurable. You can get the bindings you want. It also has configurable colour schemes, which can be looked upon

as a way to 'skin' your player.

The most exciting feature about CMus is the wide range of support for various codecs. It can play Ogg Vorbis, FLAC, MP3 (with *libmad*), WAV, AAC or WMA—i.e., almost every major format available these days. Moreover, it also supports almost every output plug-in—OSS, ALSA, *libao*, ARTS, SUN, and also WaveOut on Windows.

Using CMus

Using CMus is not very difficult if you already know or use vi(m). Let's start by adding some music to the library (yes, library—more on that later). Just type in the following command:

```
:add PATH
```

...and you will see a list of all the subdirectories and music files in the main window. If you haven't changed the keybindings yet, then, by default, you can use space to open the directory tree to list subdirectories. The files will be listed in the left side pane (refer to Figure 1). You can move between the two panes using the tab. Some more basic commands are listed below:

- `:clear` will clear the playlist
- `:save ~/playlist_name.pls` will save the current playlist into *playlist_name.pls* inside your home directory
- `:load playlist_name.pls` will load *playlist_name.pls*
- `:set softvol=true` will enable software volume control

The player supports five views and two option windows. The default window is at 1. If you press the numbers from 1 to 7, you will cycle through all the seven modes, which are:

1. **Library view** (key 1) is also the default view mode. It includes an artist/album tree, with albums sorted by year
2. **Sorted library view** (key 2) provides a simple, common playlist view with custom sorting
3. **Playlist view** (key 3) is an editable playlist
4. **Play queue view** (key 4) displays a queue of tracks
5. **File browser** (key 5) is a directory browser
6. **Filters view** (key 6) is a list of user-defined filters
7. **Settings view** (key 7) displays keybindings and options

As in the case of vi(m), you can start searching by using the / (slash) operator; pressing *N* brings the next search result. You can change the keybindings according to your preference. Copy the *rc* file from */usr/share/cmus/* to *~/cmus/*. Make the changes according to your preference to get the keybindings you want. On the other hand, you can achieve the same from the view 7 (Settings view). Just hit *Enter* after highlighting the option you want to change. It will appear in the command line. You can then make the changes and hit *Enter* to save them.

The library in the player, which you can see in the Settings view at the bottom of the page, allows you to either play the entire library (i.e., the music files you added using `:add`) or the playlist you have. Once added, the music tracks will be persistent and are only removed

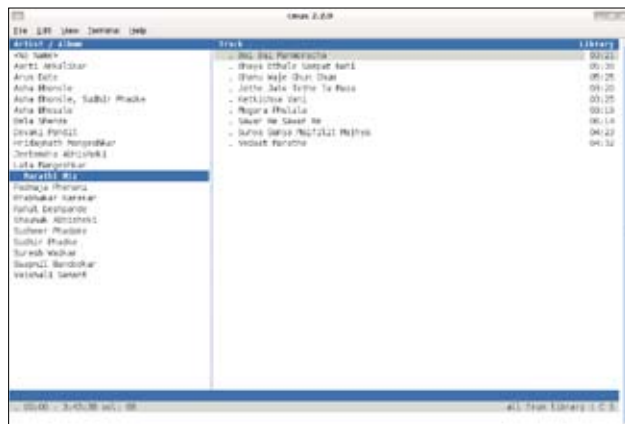


Figure 1: The CMus music player



Figure 2: Music library

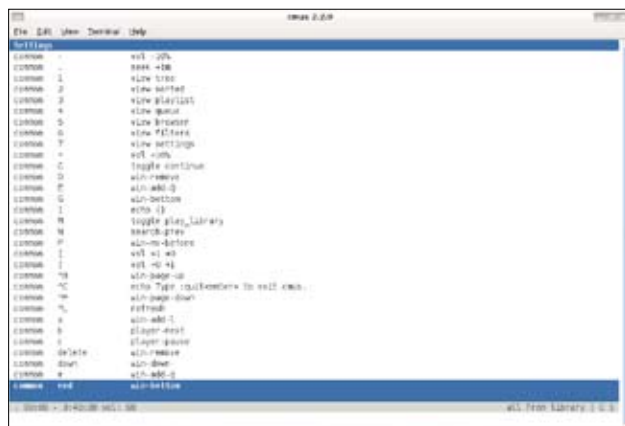


Figure 3: Key bindings

when you select them (*Space* in view 1/2) and press *Delete*. Once you have selected the songs, you can press *y* to add the songs to the playlist or *e/E* to add it to the end/start of the queue. You can also save/load the playlists as mentioned above.

Colouring

CMus also allows users to change the colour of the player windows and borders. In view 7 (the Settings view), you can view the current settings for the colour scheme. If you want to change a value, make your

changes in the `rc` file or directly from the Settings view as in the case of keybindings.

Shortcomings

As already said, being a console application, this player lacks obvious features such as cover flow, visualisations and other fancy plug-ins. Two major issues for me are that there is no option to fetch the lyrics of the currently playing song and no simple way to submit the tracks to the audio scrobbling service such as last.fm or librefm. The addition of these two features would make CMus the first choice for me. Anyway, these features have been requested and the developers have said that they will be added in the next release.

Last.fm support

Since I feel it's a must to show off my music status to everyone who cares, I scrobble the tracks I listen to in last.fm. Obviously, it was one of the first things I tried to achieve by whatever means possible. As I said, there is no 'easy' way to integrate the last.fm scrobble support. There is, however, a hard way to achieve this. Don't worry; you don't have to write code. You only need to compile the latest git version from the source. Developer Frank Terbek wrote a patch to achieve this. Get the current copy of the git repository from:

```
git clone git://repo.or.cz/cmusb.git
```

Once you have that, download the patch from http://www.hci-matters.com/blog/wp-content/uploads/2008/05/06/cmusb_audioscrobberBETA41-githead.diff. Now place this in the source directory of CMus and patch it:

```
patch < cmusb_audioscrobberBETA41-githead.diff
```

Once this is done, use the usual procedure to install a software from source (`./configure && make && make-install`). Then, open the CMus rc file (if you don't have one, make it under `~/cmusb/rc`) and add the following lines to it:

```
set as_enable=true
set as_user=(Your Last.fm Username)
```

```
set as_pass=(Your Last.fm Password)
```


Insert your user name and password at the appropriate places and you are done! Now the CMus player will scrobble the tracks you listen to at your last.fm account. To turn the scrobbling off, just change the first option value to false.

Remote controlling

If you are running CMus on a different workspace, or just want to change the song from an IRC client, you can use a program called `cmusb-remote` to change the state of the player. Apart from the usual features of play, pause, next and previous, you can also load new playlists or empty the queue. It's a very handy tool to create scripts for standalone console (or other) applications. Type `cmusb-remote --help` to get the commands that the application supports.

Finally...

The major advantage of CMus is the speed. It's very lightweight and I have been running it on my old system with a 500 MHz Celeron processor and 96 MB of RAM on Debian without a single glitch. The maximum RAM used was less than ~20MB at all times. With the likes of Amarok and Exaile taking more than 50 MB of RAM, CMus is a great alternative for your low-end systems. Sure, it's not at all as good looking as the other two, but certainly as powerful... if not more so.

On the down side, with no support for lyrics fetching and artist information retrieval available for CMus, many people will be reluctant to use it. The last.fm support, though technically there, still lacks the ease and usability that the other players support. With most of the users requesting these features, the developers have agreed to include the patches in the next major release, so it looks like CMus has a very bright future ahead of it. My call? You must try it at least once to know its power.  **END**

By: Aditya Shevade

National Talent Scholar, Aditya Shevade, a final year electronics engineering student, takes keen interest in programming and electronic design. A Linux user for the past two years, he enjoys playing the keyboard and is a good photographer. To know more about him, visit www.adityashevade.com

TechnoMail - Enterprise Email Server
Anti SPAM, Anti Virus, Email Content Filtering

Firewall, Internet Access Control
Content Filtering, Site Blocking

Bandwidth Management System

Managed Email Hosting Solutions

TechnoInfotech®

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

Explore the Save Mode With FOSS



With the recession clouds only just beginning to recede, FOSS seems the best road to take in order to slash IT costs.

*H*ave you ever gone scuba diving? On surface, the ocean looks like a vast expanse of water, mighty and gorgeous.

Once under water, you discover a paradise of marine life. Free and open source software (FOSS) holds the promise of a similar paradise for companies grappling with the receding recession. At first encounter, you would fall in love with FOSS for the scalability or improvement of the IT infrastructure it promises. But dig deeper and study companies that adopted FOSS, and you'll realise there is only one known-yet-under-rated jewel this software provides: lower costs.

Rahul De, Hewlett-Packard chair professor, Indian Institute of Management, Bangalore, recently released a study on the *Economic Impact of FOSS in India*, involving a sample of 20 organisations—including large corporations, SMEs, government

organisations, educational institutes, NGOs, etc. He estimates India will save around Rs 10,000 crores in 2010 if about half the personal computers sold in India come with free software instead of the proprietary software.

"The cost savings one gets by deploying FOSS are immense. From costs incurred on acquiring and installing software, money spent on complementary software for security, software upgradation costs and switching costs associated with moving from one type of software to another, corporates moving to FOSS saved on all fronts," says De.

Life Insurance Corporation (LIC) of India saved a crore on licensing costs by running around 18,000 desktops on free software, apart from thousands of servers; and the IT@School project of Kerala saved around Rs 49 crores by placing free software on 50,000 desktops in 2800 schools across the state.

The cost factor and its importance

But is just cost reason enough for an organisation to adopt FOSS? Though an important factor, cost isn't the sole reason to adopt FOSS. Free software provides numerous benefits like greater scalability, improved performance, security, etc, which most often outweigh the cost benefit.

"Most companies experiment with FOSS because of the numerous benefits it provides, like improved performance that includes aspects of stability, interoperability, operational ease and maintenance, security, absence of a vendor lock-in, etc. Cost is an influential factor but not the sole reason for FOSS adoption. A lot of the intangible cost benefits experienced after adopting aren't even known by most people during adoption," says De.

For organisations like Sheela Foam (manufacturers of the Sleepwell brand of mattresses), Eveready Industries (one of India's leading battery manufacturers), and IT For Change (a Bangalore-based NGO), it was the complete FOSS cookie basket that made it attractive.

"Sheela Foam turned to FOSS as we wanted a cost-effective solution to host our ERP (Greatplus) system and enjoy the flexibility to deploy new applications. We initially ran our ERP system on a legacy server in a proprietary UNIX environment. Our legacy platform had limitations, thus we were unable to introduce new applications and scale effectively. To preserve our initial investment in that platform, we had to invest considerably in new hardware and licensing fees. Thus, we migrated to FOSS in April 2009 placing three of our servers on FOSS and using Red Hat Enterprise Linux," says Pertisth Mankotia, head-IT (DGM), Sheela Foam (Sleepwell).

Says Arup Choudhury, CIO, Eveready Industries, "The cost factor played a major role in deciding to go for FOSS. Besides, we also got the level of security in an x86 platform (FOSS server) that was comparable to RISC, and improved performance because of the UNIX equivalent kernel."

Gurumurthy Kasinathan, director, IT For Change (ITfC), an NGO that migrated to FOSS in early 2007 and also a participant in the study, adds: "FOSS provides benefits that are particularly important for non-profit organisations. Of course, one of the main reasons for adopting it was low cost. With no license fees to be paid, and no upgrade choices imposed by vendors, the cost is much lower with FOSS, especially in the long term. But it could also be tailored to meet our local needs and was less vulnerable to viruses. The most influential factor in ITfC adopting FOSS was the philosophy behind using

FOSS, which is that anybody has the freedom to use, study, modify and share it."

However, for some organisations like the Institute of Informatics and Communication (IIC), University of Delhi, South Campus, another participant in the study, the cost factor had little influence. "University systems have diverse technological requirements to cope with the needs of numerous disciplines. To provide the latest state-of-the-art technology to the students, you can't wait for months to procure all those unaffordable closed-source solutions. Moreover, students should be exposed to the latest software solutions at the right time. To bring in this flexibility, we adopted an open source framework in 1998," says Sanjeev Singh, placement and project coordinator, Institute of Informatics and Communication, University of Delhi, South Campus. The institute has around 30 servers running on different flavours of Linux like CentOS, openSUSE, Debian, etc, and 100 desktops using various Linux distros.

“From costs incurred on acquiring and installing software, money spent on complementary software for security, software upgradation costs and switching costs associated with moving from one type of software to another, corporates moving to FOSS save on all fronts.

—Rahul De, Hewlett-Packard chair professor, IIM, Bangalore

Economic impact of FOSS

Licensing, service, distribution, maintenance, integration, migration, upgradation and exit costs are some areas on which an organisation saves. The economic impact of FOSS can be measured by three principal means: FOSS as a substitute for more expensive desktop operating systems and office productivity applications; FOSS as a substitute for more expensive server software; and FOSS-enabled cost savings from complementary products such as anti-virus software required on Windows desktops.

By using FOSS as a desktop operating system, one can save on the licence cost of the office productivity tools and the operating system. The price of operating systems like Windows ranges from Rs 2,000 to Rs 13,750, while the office productivity products like Microsoft Office cost anywhere between Rs 3,025 and Rs 26,800.

Elaborating on the cost savings, De says: "Most users of office products are unsophisticated, using only

a fraction of the total features provided in professional products like Microsoft Office, while a free office suite like OpenOffice.org made more sense and thus, they shifted.”

Those at Sheela Foam, that’s been using Red Hat Enterprise Linux since April 2009, couldn’t agree more. “We invested only about Rs 8 lakhs and expect to incur a saving of Rs 50 lakhs spread over three years, because of migrating from proprietary software,” says Mankotia.

A small NGO like IT For Change, running about 21 laptops and desktops on Ubuntu, and using OpenOffice.org Writer and Calc, admits to obvious cost benefits by switching to FOSS for generic applications. “FOSS has helped us save on license costs, service costs, and upgrade costs. We have saved on 21 Windows and MS Office licences which adds up to approximately Rs 3 lakhs,” says Kasinathan. Alongside, De’s study shows that the IT@schools project that replaced the Windows OS on 50,000 desktop computers in 2,800 schools across the state with a FOSS OS, saved around Rs 490 million, while an e-commerce arm of a Mumbai-based retail chain saved Rs 3 million by replacing MS Office with OpenOffice.org.

However, a lot of doubts have been raised about data migration and integration when it comes to open source applications, and regarding open source as an operating system. “We initially thought that migration would be difficult. However, we experienced that it was quite simple. We did not require much training and there were no issues of integration, scalability, etc,” says Mankotia, nullifying any doubts about transitional hiccups.

FOSS replacement on servers either as operating systems hosting application products, or as enterprise class products, too, has helped organisations enter the ‘save mode’. “We used Red Hat advanced server 4.0 and enjoyed immense savings on hardware costs, since a RISC server cost us 40 per cent more than an x86 server,” says Choudhury. New India Assurance that migrated about 1500 servers to open source and also a large part of its desktops to Linux, saved about 35 to 40 per cent on IT expenditure, which amounted to Rs 800 million, reveals De’s study. Further, Sheela Foam that shifted about three servers to open source said that its maintenance costs dropped to one-sixth after migrating to a FOSS-based system.

A Computer Security Institute/FBI report in 2006 stated that the loss from virus-related problems for about 300 firms in the US was close to \$ 50 million. Indian firms too have reported massive losses owing to attacks from viruses and worms. One estimate places the recovery time to be, at an average, 29 hours from each serious virus attack. In most businesses where time is money, this amounts to huge losses. Using a FOSS operating system on the desktop very often requires no anti-virus software. “One of the servers is used for the security firewall, proxy, etc. We have not purchased any anti-virus software, so it is a 100 per cent saving on the security

front because of using FOSS,” testifies Singh. Other perks availed by organisations going in for FOSS include a 30 per cent cut in energy bills for Sheela Foam.


Further, FOSS has been crucial to organisations pressing the innovation keys to improve the way they function. “FOSS has immense innovation potential. It provides a platform for the creative development of software and aggregation of applications. It enables organisations and individuals to experiment with, tinker with and play around with different combinations of hardware, software and applications. This has a potential to create immense benefits for the firm,” says De.

At IIC, FOSS products enable students to experiment with and hence learn about software. “OSS and innovation go hand in hand. Students choose to work on open source projects, because they get enough freedom to explore, which is a crucial component of any learning process. Now we are at a stage where we have enough projects developed over training periods. These can be used for the academic requirements of any institution. They include developments like algorithms for the Gnome rice project (for sequence mapping and motif identification), network monitoring, alumni database systems, academic portal frameworks, etc, which we release as an open source software resource toolkit,” says Singh.

Kasinathan says: “FOSS has helped foster innovation—perhaps not as much for our organisational performance as much as for the projects we undertake. We support and encourage other organisations to innovate in other areas such as education. For instance, we are collaborating with a partner organisation to create a customised distribution of Debian, including educational packages and local language support, which will be used in government schools for computer education.”

“If we have to innovate or test anything now, we can replicate our data in any machine, as our OS is not machine-dependent, which in turn has helped us innovate more. We have not only saved costs, our boundaries too have been eliminated to a great extent,” adds Mankotia. According to De’s study, others like New India Assurance and the IT@Schools project too have innovated successfully using FOSS.

In a nutshell

Altogether, FOSS has been instrumental in helping organisations achieve superior benchmarks without huge costs. People have begun to appreciate the benefits of using FOSS to create, connect and share. And once begun, this trend will only grow. **END** 

By: Vanisha Joseph

The author loves to experiment and writing for LINUX For You is her latest experiment. So, beware! Just a minute, she also happens to be a journalist during the day.

Getting Started with GNU Octave



Part 1

A lot of what GNU Octave provides revolves around manipulating matrices, so let's begin our journey with them.

GNU Octave is a software tool for numerical computations. Octave has extensive tools to solve common numerical linear algebra problems, find the roots of non-linear equations, integrate ordinary functions, manipulate polynomials, and integrate ordinary differential and differential-algebraic equations. Octave is available for use on UNIX systems (Linux, Solaris and Mac OSX) and Windows.

If your study or play involves numerical computations, Octave is for you. It has been designed with MATLAB compatibility in mind, so with careful design you could write scripts that run on both MATLAB and Octave. It is easily extensible and customisable via user-defined functions written in Octave's own language, or using dynamically loaded modules written in C++, C, Fortran, or other languages. As we go ahead, we

will look at how we can carry out these functions and more, in Octave. For now, let's start using Octave. The first couple of articles in this series will deal with using Octave for matrix manipulation and linear algebra, in general. I assume readers have some basic knowledge of the mathematics of matrices and linear algebra, and I will not make any effort to explain that.

Installation

As always, you have two options: to use the package manager of your Linux distribution, or compile it from sources. It is always easier to do an *apt-get install* or its other cousins. Please find the name of the Octave package for your own distro.

Let's play with Octave

Fire up the terminal and type `$ octave` and you should see something like the following:

GNU Octave, version 3.2.2

Copyright (C) 2009 John W. Eaton and others.

This is free software; see the source code
for copying conditions.

.
.

...and finally the Octave prompt:

```
octave:1>
```

So, we are using Octave 3.2.2. You can *quit* or *exit* by typing in the same—not that I want you to quit; we have only just started. A lot of what Octave provides us revolves around manipulating matrices, so let's start off.

Define a simple 2x2 matrix, X:

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \quad (1)$$

This is how we can define it in Octave:

```
octave:1> x= [1,2,3,4] x = 1 2 3 4
```

...and let's define another, Y:

$$\begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix} \quad (2)$$

```
octave:1> Y= [5,6,7,8] Y = 5 6 7 8
```

We have two square matrices, X and Y. What shall we do? Let's find the product:

```
octave:13> X*Y ans = 19 22 43 50
```

Easy, isn't it? We defined two matrices, X and Y, and then found the product by simply X*Y. It's like how you would multiply two integers or real numbers. You get the idea, right? The individual columns are separate by a ',' and the rows by a ';'. Let's move on.

Things to try out

We have X and Y defined now. Let's try something like, Z=[X,Y] in Octave. What do you see? You should see the X matrix augmented or combined with the Y matrix:

$$Z = \begin{bmatrix} 1 & 2 & 5 & 6 \\ 3 & 4 & 7 & 8 \end{bmatrix}$$

We have so far had only real numbers in our matrices. How about complex numbers? What does Octave have to say about them? Let's see.

```
octave:18> XC = [1+2i,3;3+2i,4]
XC =
```

```
1 + 2i   3 + 0i
3 + 2i   4 + 0i
```

```
octave:19> YC = [1+2i,3;3+2i,4]
YC =
```

```
1 + 2i   3 + 0i
3 + 2i   4 + 0i
```

```
octave:20> XC*YC
ans =
```

```
6 + 10i   15 + 6i
11 + 16i   25 + 6i
```

Great! Octave handles them well too. As you can see, when you define some elements of a matrix as complex, the real elements are also written in complex form.

Now that we have had a taste of Octave, let us try to understand some basics, and after that we shall resume our fun with matrices on a more serious note.

Built-in data objects

We have already used numeric data objects in our matrices. Besides numeric data objects, we also have string data and data structure objects.

- **Numeric data objects:** Octave's built-in numeric objects include real, complex and integer scalars and matrices. All built-in floating point numeric data is currently stored as double precision numbers. For the exact values for the maximum and minimum possible real numbers that can be represented on your system, type *realmax* and *realmin*, respectively.

- **String data objects:** A character string in Octave consists of a sequence of characters enclosed in either double or single quotations. Strings won't be useful to us for the purpose of this article, so I won't talk about them.
- **Structure objects:** You can also have C style structures in Octave. You can have a number, a matrix and a string, and combine them under a common structure object. For example:

```
octave:4> x.a=4
x =
{
a = 4
}
```

```
octave:5> x.a=4;
octave:6> x.b='foo baz'
x =
{
a = 4
b = foo baz
}
```

```
octave:7> x.b='foo baz';
octave:8> x
x =
{
a = 4
b = foo baz
}
```

Numeric data types

We can use the numeric objects to create higher level numeric data types like matrices, ranges and cell arrays. We've worked a bit with matrices and we will do a lot more with them now. I shall introduce the others later on in this series of series.

Let's now have some more matrix fun, shall we?

Determinants, inverses and singularity

For anyone who has taken a basic linear algebra course, close on the heels of matrices follows the concept of determinants. Let's try this out in Octave:

```
octave:1> X=[2,2,3,4]
X =
```

```
2 2
3 4
```

```
octave:2> det(X)
ans = 2
octave:3> X=[2,2,3;3,4,3]
X =
```

```
2 2 3
3 4 3
```

```
octave:4> det(X)
error: det: argument must be a square matrix
octave:4>
```

Linear algebra tip: A determinant is defined only on square matrices. If we view a square matrix X , as a linear transformation on a vector A , the determinant of X determines the scale factor of the transformed vector A . That is, if A encloses a finite area in space, the transformed matrix AX will enclose double the area, assuming $\det(X) = 2$.

By any chance, would you remember the case of $\det(X) = 0$? Such a matrix X , is said to be singular. Singular matrices are not invertible. This means that we cannot have an inverse of such a matrix. Let's see what Octave says:

```
octave:8> X=[2,2,2,3]
X =
```

```
2 2
2 3
```

```
octave:9> inv(X)
ans =
```

```
1.5000 -1.0000
-1.0000 1.0000
```

```
octave:10> X=[2,2,2,2]
X =
```

```
2 2
2 2
```

```
octave:11> inv(X)
warning: inverse: matrix singular to machine
```

```
precision, rcond = 0
ans =
```

```
Inf Inf
Inf Inf
```

Utility matrices

Octave also has functions to generate some utility matrices, such as an identity matrix using the command *eye*, a zero-filled matrix using *zeros*, and others. Please consult the Octave software documentation for the detailed treatment of these functions.

The Hilbert matrix

The Hilbert matrix, named after the great German mathematician David Hilbert, is a square matrix with its entries being of the form: $H(i, j) = 1 / (i + j - 1)$. Octave has a built-in function *hilb* to generate Hilbert matrices. For example:

```
octave:16> H= hilb(3)
H =

1.00000 0.50000 0.33333
0.50000 0.33333 0.25000
0.33333 0.25000 0.20000
```

The above is a Hilbert matrix of order 3. Now calculate the determinant of H :

```
octave:18> det(H)
ans = 4.6296e-04
```

Now, calculate $1/\det(H)$:

```
octave:19> 1/det(H)
```

```
ans = 2160.0
```

This is a curious fact. The determinant of a Hilbert matrix is always the reciprocal of an integer. (As stated in the online encyclopaedia of Integer Sequences at www.research.att.com/njas/sequences/A005249)

If you liked that famous matrix, Octave has a few more. Refer to the documentation to play around with them.

Next up

We covered a lot of matrices in the first part of this series on GNU Octave. We will work with more, focusing on solving linear equations, in the next part of this series, after which we will move to other areas of GNU Octave. A few things to explore on your own, besides what we have already talked about, could be:

- Command-line editing in Octave
- Playing with the Octave prompt
- Experimenting with some of your legacy MATLAB scripts in Octave

Happy exploring with GNU Octave and feel free to write in with your queries and suggestions.



Resources

- GNU Octave: www.gnu.org/software/octave
- GNU Octave Function reference: octave.sourceforge.net/doc/index.html
- Wikipedia articles on Determinants, Matrices and Hilbert Matrix

Acknowledgements

Considering the heavy usage of mathematics formulas, LaTeX was an automatic choice to write this article and thanks to the The Not So Short Introduction to LaTeX [<http://ctan.tug.org/tex-archive/info/lshort/english/lshort.pdf>] for insights on how to write mathematics well in LaTeX. Hevea [<http://pauillac.inria.fr/hevea>] was used for LaTeX-to-HTML conversion and it is awesome.

By: Amit Saha

The author blogs at amitsaha.wordpress.com and he welcomes queries and suggestions at amitsaha.DOT.in@gmail.com.



Think Parallel with OpenMP

Without stepping into your C or C++ based application's algorithmic design, you can still add threads to your existing code to achieve performance upgradation on a multicore processor. How, you may ask. This article tells you!

Do you like your existing sequential C or C++ code multi-threaded so you can run the performance-sensitive codes on multicore processors? And do you want all this with minimal alteration to your code?

Hmm... well, have you ever heard of OpenMP? Let's look at how it can help you.

OpenMP (Open Multi-Processing) is an open standard for writing shared memory parallel programming in C, C++ and Fortran, in the form of mainly compiler directives, library routines and environment variables. This open standard is supported by multiple architectures/platforms including UNIX/Linux and Windows NT. The Architecture Review Board (<http://www.openmp.org>) that owns and maintains the OpenMP brand is a non-profit corporation that has AMD, Cray, Fujitsu, HP, IBM, Intel, NEC, The Portland Group Inc., SGI, Sun Microsystems and Microsoft as its members.

I assume readers are familiar with parallel

programming. As a developer working in C language, my obvious choice of language to show you OpenMP program examples is C, but with little (or rather, no) tweaking, any C++ programmer can also use them. GCC 4.3.2 or higher is required (a whole list of supported compilers can be found at <http://openmp.org/wp/openmp-compilers>). You can use any compiler that supports OpenMP; however, I will use GCC for all my examples in the article.

The OpenMP programming model

OpenMP follows the fork-join parallel execution model. An OpenMP program starts its execution as a single thread (called initial thread) and assumes the execution sequentially. This sequential program gradually evolves as a parallel program, as a 'team of threads' is created when the master (or initial) thread comes across a *parallel* construct in the code to share the workload. Each task is assigned and tied to a different

thread in the team. At the end of the parallel construct, all threads synchronise and get terminated. Only the master thread continues the further execution. Figure 1 depicts the fork-join model of OpenMP.

Within a single program, you can mark any number of parallel constructs.



Note: Team of Thread = Master + Workers.

When any thread team encounters a work sharing construct, the work inside the construct is offloaded and distributed among its members, and gets executed cooperatively, instead of the whole work being executed by every thread, multiple times. Each thread just has its portion of work, which it executes synchronously.

When can you consider OpenMP?

- When data parallelism is required. According to Wikipedia, data parallelism (also called loop-parallelism) is achieved when each processor performs the same task on different pieces of distributed data.
- When a block of code is required to be executed in-parallel.
- Importantly, when you don't want to bring modification to your existing application design, and yet want to achieve parallelism. All you need to do is to carefully analyse your code and insert appropriate compiler directives, and compile your code with the compiler that supports OpenMP. With compilers that don't support OpenMP, just ignore these directives.

How is OpenMP used?

Well, as already mentioned, an OpenMP program consists of: compiler directives, library routines and environment variables. OpenMP directives for C/C++ are specified with the pragma pre-processing directive, and each directive starts with `#pragma omp`.

You need to first specify how many threads you want to create for your parallel regions. To do so, follow the steps given below:

- At run-time with `omp_set_num_threads`:

```
#include <omp.h>
omp_set_num_threads(int num_threads)
```
- Or, via environment variable `OMP_NUM_THREADS`:

```
export OMP_NUM_THREADS=<specify the number of threads you want>
```

Function prototypes and types are in the `omp.h` file.

Finally, you can compile your OpenMP code with the `-fopenmp` switch:

```
$ gcc <your C file> -o <output file> -fopenmp
```

The next section deals with useful OpenMP directives




Figure 1: OpenMP's fork-join model

for *parallel* processing.

Some basic OpenMP constructs

The fundamental construct of OpenMP is *parallel*. This construct starts *parallel* execution. Without this construct, your code will remain sequential:

```
/* Sequential Code here...(Master Thread) */
#pragma omp parallel
{ /* A team of N Threads starts here */
  /* Team of Threads = Master + Worker Threads */
  /* Parallel computing here...*/
}
/* Master Thread resumes back here */
```

Just a *parallel* region alone doesn't specify how the work will be shared. And thus all threads will end up redundantly executing all the work—that is, no work sharing and, therefore, no actual speed gain. Compile the example given next with GCC and the `-fopenmp` switch, and run it to see the output.

```
#include <omp.h>
#include <stdio.h>

int main (int argc, char *argv[])
{
  int nthreads, tid;

  omp_set_num_threads(4);
  /* Just fork a team of threads */
  #pragma omp parallel
  {
    /* Obtain thread number */
    tid = omp_get_thread_num();
    printf("Hello World from thread = %d\n", tid);

    /* Only master thread does this */
    if (tid == 0)
```

```
{
nthreads = omp_get_num_threads();
printf("Number of threads = %d\n", nthreads);
}

} /* All threads join master thread and disband */

}
```

With work-sharing constructs, OpenMP splits the work of parallel regions and distributes them among threads. OpenMP work-sharing constructs are:

- *loop* (for a loop in C/C++)
- *section*
- *single*

Parallelising loops are the most important constructs in OpenMP. In your code, you need to find the most time-consuming loops, and you can split the task inside it among the team of threads.

See the following snippet of code:

```
int a[5000], i;
#pragma omp parallel for
for(i=0;i<5000; i++)
{
    a[i] = i;
}
```

The above snippet demonstrates how to parallelise a simple loop using the loop construct. The loop iteration value (variable 'i') is private, by default, to each thread in the team, and each thread is assigned a fixed-size chunk. It is important to note that the loop in the work-sharing region must be iteration-independent, so that each thread can safely execute its chunk independently -- in any order without carrying any previous loop-dependency.

Second in work-sharing constructs, is *section*. With the *section* work-sharing directive, individual blocks of code are distributed among the threads in the team. Each thread can execute one or more sections and each section is executed exactly once. Some threads may even remain idle if the sections are too few. The following code snippet then gives an imprint of the section construct:

```
#pragma omp parallel sections
{
    #pragma omp section
    {
        x_cal();
    }
    #pragma omp section
    {
        y_cal();
    }
    #pragma omp section
    {

```

```
z_cal();
}
}
```

The third in work-sharing constructs is *single*. The *single* construct denotes a block of code that is executed by any one thread from the pool of threads, including the master thread. The other threads in the team, which do not execute the block, simply wait at the implicit barrier at the end of *single* construct, unless a *nowait* clause is specified. One possible use of this construct is to initialise shared variables.

```
#pragma omp single
{
    a=10; /* N-1 threads bypass this block */
}
```


How to do thread synchronisation

Well, the important thread synchronisation constructs are:

- **Master directive:** The block that is executed by only the master thread.
- **Critical directive:** The block that is executed by only a *single* thread when taking care of the race-condition.
- **Barrier directive:** An explicit barrier at a point at which the construct appears.
- **Taskwait directive:** Specifies a wait on the completion of the child task (generated at the beginning of the current task).
- **Atomic directive:** Ensures that a specific storage location is updated atomically.

OpenMP also provides a set of general-purpose, run-time library routines that can be used for synchronisation. These general-purpose routines operate on OpenMP lock variables. Some of these important routines are, *omp_init_lock()*, *omp_destroy_lock()*, *omp_set_lock()*, and *omp_unset_lock()*.

So, what's next?

We've discussed OpenMP here with the purpose of sharing its basic idea. You can further consult "OpenMP Application Program Interface", which you can download from the OpenMP official site (www.openmp.org).  **END**

References

- OpenMP Application Program Interface Version 3: www.openmp.org/mp-documents/spec30.pdf
- OpenMP tutorial: computing.llnl.gov/tutorials/openMP
- Wikipedia on OpenMP: en.wikipedia.org/wiki/OpenMP

By: Vivek Kumar Prasad

The author is a system engineer at IBM India, and works for IBM-ISS (Internet Security Systems). He is based in Pune and you can reach him at mail_DOT_vivekkumar_AT_rediffmail_DOT_com.

An Exclusive B2B Exhibition - With a Senior Level Conference



5th Information & Communications
Technology Exhibition & Conference

8 - 10 May 2010

Karachi Expo Centre



- An exclusive B2B event showcasing the latest IT and Communications Technologies
- The 5th High level Industry Conference being held in concurrence with the exhibition
- Presence of business delegates, industry professionals and stakeholders
- Top level networking opportunities for the participants through one-to-one meetings, technical sessions and product workshops

www.connectit.com.pk

To participate, please contact:



Pegasus Consultancy (Pvt) Ltd

2nd Floor, Business Centre,
Mumtaz Hassan Road, Karachi - Pakistan
Tel : +(92 21) 111 734 266 (PEGCON)
Fax : +(92 21) 3241 0723

E-mail : info@connectit.com.pk URL : www.pegasus.com.pk

Supported by:



Engineering
Development Board



Board of Investment
Government of Pakistan



Pakistan Software
Export Board

E-Media partner:





Victory: The Age of Racing by Vae Vivrics [www.vaevivrics.it/projects.html] uses OGRE for rendering.

Getting Started with GAME Development

Creating a game is like creating a movie. It involves money, and the time and effort of a lot of people—or so I thought. Then I bumped into independent developers who were making games on shoestring budgets, yet keeping abreast with bigger production houses that were backed by million-dollar budgets.

Game development doesn't necessarily mean that you have to know programming, but this article focuses on development rather than anything else. Graphics programming is different from your regular application development. Regular applications are event-driven. Games are not.

Event-driven applications are almost all applications that you use every other day, like a calculator or a spreadsheet program—they require inputs from a user to move along.

When I started working with Visual Basic a long time ago, I didn't really know anything about how games worked. However, I was pretty used to the idea of event-driven applications.

Higher-level languages like Visual Basic take care of a lot of low-level work using intermediate runtimes. These leave you with a bit of reduced flexibility, unlike when dealing with the core platform itself.

Even later, after switching over to C++,

graphics programming still remained an elusive dream that left me confused because none of the applications I wrote behaved like games—they rarely 'acted' for that matter, and then came the game loop.

The game loop

The game loop is the heart of a game. It's a regular loop that continuously executes commands that keep the whole game world, as well as its occupants, running.

There were only ideas in the beginning. To move something on screen, I could use a loop or a keyboard input event—yes, even that was good. But my program flow was stuck in the loop of that specific event. If a car was being moved by my arrow keys, nothing else worked while that happened. To make my whole game world live and make everyone happy, the 'pros' of the black art of graphics programming came up with something called the game loop.

Regular applications waited around for

input, while our game loop looked something like this:

```
bool bGameRunning; // variable to hold a boolean if the game is running or not

int main() {
    initWindow(); // we create and initialize a window
    initGraphics(); // we initialize the graphics interface
    initSound(); // inits sound interface with the sound hardware
    initInput(); // initialize input
    initPhysics // initialize physics
    initNetwork(); // initialize network
    bGameRunning = startGame(); // starts game returns true if all went well

    while(bGameRunning) { // this is our loop, while the game is running keep
        looping
        updateInput(); // updates the Input
        checkCollision(); // check for collision player position + world + NPC
        updateAI(); // update all AI/NPC in game
        updatePlayer(); // update player according to input
        updatePhysics(); // update Physics of the game world
        updateServer(); // send messages to server + other players
        updateGame(); // update game, see if it has ended or not
        updateDisplay(); // update screen according to all changes
    }

    return 0; // game ended, program flow exits the loop and we exit the program.
}
```

A constantly updating chain of commands keeps the game world and its creatures alive even when there is no update from the user. As you can imagine, more functions and functionality would be added to the list for handling things like networking, audio, sound, music, more input devices, physics, etc. As this system grows, you should bring your object-oriented programming skills to the table to correctly classify the objects. So you may have an idea by now of what game development might involve. There's a heck of a lot more, however!

What tools and technology you use largely depends on how you want to go on developing the game. Game development involves a whole lot of tools and technologies—from your content-creation pipeline to developing tools for your game and, finally, to the actual game itself.

A great deal of assets for the game are created using software like Autodesk's Maya and 3D Max. Development usually takes place using very expensive IDEs, and management software like Microsoft's Visual Studio and the likes. If you are not making middleware (software and technologies that facilitate the job of a middle man in game development --which could range from file formats, 3D graphics, physics, network engines, to whole game engines), then you would be using them and they don't come cheap. Apart from indie licences of some game engines (Torque, to name one), almost all middleware is expensive and comes with very restrictive licences.

Enter, free and open source software!

You can save a lot of time and money by carefully selecting a set of FOSS tools in your game development lifecycle. Apart from saving time and money, FOSS tools and technologies have been time-tested and moulded into what they are by strong communities of developers, as well. There are a plethora of tools and technologies out there to help you get started and get along. So let's look at some of them.

Graphics engines and general-purpose libraries

Graphics engines are libraries that interface with your graphics hardware and draw everything you see in your game. You can start writing this component from scratch, focusing on proprietary DirectX or free OpenGL—both industry standard interfaces to graphics hardware.

Writing one would involve selecting an interface; what scene graph algorithms to use; how to load and manage 3D data, and how to send it to the graphics hardware; creating the means to set up light; working with textures; animation; IK; and a whole lot more. You can save a lot of trouble and jump into developing your game rather than spending time deciding how you do translations to matrices, by selecting a ready-made engine that does all this for you, out-of-the-box.

There are many free and open graphics engines out there—OGRE, Irrlicht, Crystal Space, Panda3D and the Blender Game Engine are some that stand out from the rest. Blender is primarily a 3D modelling and animation software. It can also do post-production work like editing and compositing. Apart from serving your content-creation pipeline, it has a powerful game engine as well. Some of these graphics engines sometimes provide you with not only graphics, but also other functionality like the GUI, input, post effects, file management, 3D functions, mathematics, etc.

Panda3D from CMU and Disney gives you a whole lot of tools that you can use for all your game development needs. Some of them are code for the scene graph and for rendering, while some are for graphic asset loading, timers, tasks and event handling, sound, input, video, the GUI, post FX, procedural content generation, collisions and physics. These tools also include code for mathematics, finite state machines, scripting, networking, a huge set of sample programs and code to start from, packaging and distribution solutions, debugging and performance tuning, and tools for third-party software.

The list of libraries can go on and each has its pros and cons. Testing and selecting the libraries that you require solely depends on the kind of game you are developing and the requirements you have. There are a lot of general-purpose libraries like the infamous SDL, ClanLib, Allegro and SFML that do a lot more than talk with your graphics hardware, providing file system functions, file management, mathematics, 3D functions, GUI, input and networking, to name a few. There are more libraries like the OpenSceneGraph, OpenGLut and FreeGlut that deliver smaller frameworks for your 3D application needs.

2D development

There are some libraries that you might want to make use of for your game development needs. A few, like Allegro, ClanLib and SFML are good, but the learning curve might seem a bit steeper.

Helper functionality: image, video, sound, networking

Image loading and image functions are sometimes required, again depending on the nature of your game and application. There is the infamous DevIL (Dev Image Library) or you can try your hand with FreeImage or libPNG. Most of the graphics engines have built-in capabilities to load, create and manipulate images at a very basic level.

Some graphics engines provide functionality to decode video and render them to textures out-of-the-box. But for the others, you can always use libraries like FFmpeg, XviD or Theora.

To rock your audience with the right amount of music, the best option is to use a library like Audiere. OpenAL also provides you with solid methods to access your sound hardware for your game or application with minimal hassles. libMikMod and Ogg Vorbis are also worth trying out. BASS and FMOD (latter being more optimized) have to be mentioned since they are awesome too (free for non-commercial use only). irrKlang is also another 'free for non-commercial use' library like FMOD and BASS. Then again, if you are lucky, your 3D engine might already have a wrapper for some audio library, saving you the trouble of rolling out your own.


Networking is an option you might want if you are going for a multi-player game. RakNet or OpenTNL are easy-to-use high-level libraries especially developed with games in mind. You can also try out the networking options present in SFML, if you are game. There are low-level alternatives that are a bit harder to understand and implement, like the Boost ASIO, HawkNL, etc.

The extra mile

There is always something more to add. If you are coding up an editor for a game, a great addition to it would be a scripting element—Lua being the most popular around, with scripting languages like Python, Ruby, AngelScript and even JavaScript following right behind. Squirrel, GameMonkey and even ChaiScript (www.chaiscript.com) are popular options too.

For your games to react and work like the real world, you need real-world physics. Just like we picked up every other library, we get to choose from some wonderful physics libraries that simulate real world physics. Two libraries that you must look into are the Open Dynamics Engine (ODE) and Newton Game Dynamics. There is also Tokamak and the now-commercial Novodex (a.k.a PhysX) to check out. Some really brilliant games like Gish and Crayon physics would not have been possible without Physics in 2D. To do just that there are some awesome libraries out there like Chipmunk Game Dynamics (code.google.com/p/chipmunk-physics) or one of the many ports of the APE (www.cove.org/ape).

To add a bit more life to your NPC, you can always turn to AI libraries like OpenSteer and MicroPather. FEAR is also something you might want to look at.

You can always use more tools in your toolkit and if they're FOSS tools, that's even better. Boost libraries are something you can always rely on when providing random numbers or giving you smart pointers. wxWidgets (wxwidgets.org) can provide you with an easy interface and GUI for your game editor in no time. You can use Doxygen to document the functions you make, and NSIS to pack your game and send it to all the gamers around. With so many awesome tools and technologies in your arsenal, your game development needs are pretty much covered. All that is left is to go develop that game!  **END**

Acknowledgements

- Links to all libraries mentioned here can be found at www.twilightsebrace.com/personal/gamelibs.php. The list is maintained by Ben Sizer. A similar list can also be found at www.ebonyfortress.com/blog/free-game-development-libraries.
- The Free Resource Thread in the *gamedev.net* contest forums can also be helpful to get you started—www.gamedev.net/community/forums/topic.asp?topic_id=324643.
- If you are an indie or if you just like making games like the rest of us here, do drop by at indiegamedev.in or follow us @ingdin on Twitter. You can also find us on Facebook at www.facebook.com/pages/Indie-Game-Development-India-Community/101358037519.
- All images are from the Ogre Showcase Forums and Ogre Flickr Gallery. A word about Ogre: OGRE (Object-Oriented Graphics Rendering Engine) is a scene-oriented, flexible 3D engine written in C++, designed to make it easier and more intuitive for developers to produce applications utilising hardware-accelerated 3D graphics. The class library abstracts all the details of using the underlying system libraries like Direct3D and OpenGL, and provides an interface based on world objects and other intuitive classes. Find out more about it at www.ogre3d.org.
- Ogre Gallery: www.ogre3d.org/gallery
- Ogre Showcase Forums: www.ogre3d.org/forums/viewforum.php?f=11&sid=3a4c8230b5decd5d06b384a15c997d97

By: Yadu Rajiv

The author is a UI/UX/designer at OSSCube and an independent game developer. He is also the co-founder of the Indie Game Developers India (InGDIn) Community [indiegamedev.in]. You can read his rants at chronosign.com/rant and follow him on Twitter at @yadurajiv.



Parsing Complex Declarations in C

C is infamous for its complex (and clumsy?) declaration syntax. In this article, we'll look at a simple technique to parse such complex declarations in your head.

C (and of course, C++) is infamous for its complicated declaration syntax that invariably causes beginners to stumble. Complex declarations are hard to decode and understand. However, this is an essential area to master because you require it in real programming applications.

There is a simple traditional technique available called the 'clockwise' rule, which states, "Take any declaration, start with the innermost parenthesis (in the absence of parenthesis, start with the declared item's name), and work clockwise through the declaration going to the right first." For instance, with...

```
void (*arr[10]) ( );
```

...as the rule says, start from the innermost parenthesis and take only `(*arr[10])` to begin with. In that part, start with the variable name, which is 'arr'. Work clockwise and read it. "arr" is an array[10]. We hit ')', so go back and read *, and so the sentence reads "arr" is an array[10] of pointer".

Then continue reading the remaining parts: "...to function returning type void". Reading it fully, it comes out that: "arr" is an array[10] of pointer to function returning type void". If you get confused about when to read left and when to go right, you can use this guideline: "Go to the right when you can, and to the left if you must". Simple, isn't it?

Just as an aside, this rule works for all declarations of original K&R C. But a slight change is required to accommodate the ANSI C addition of qualifiers like `const`, `volatile` and `restrict`. If they are used, read it together with the closest entity in the right-hand side. For example:

```
const int* (*foo)(void);
```

With this modified clockwise rule, we read it as: "foo is a pointer to function taking no arguments and returns a pointer to const int". Looks fine, right?

Enough with dummy examples; let's look at a real (and well-known) example. This declaration is from `<signal.h>`:


```
void (*signal(int sig, void (*func)(int)))(int);
```

Shall we parse it? "signal" is a pointer to the function taking two arguments 'sig' of int type and 'func', which is a pointer to a function taking int as an argument and returns void type; it returns a pointer to the function taking int as argument and returning void!".

Let's accept that this is tough (even with all the help from the simple 'clockwise rule')! *typedef*'s can simplify complex declarations (see the Dec'08 JoP column). We'll use *typedef* to simplify this declaration:

```
typedef void (*sig_t) (int);  
sig_t signal(int sig, sig_t func);
```

Yes, that's certainly more readable, isn't it? In general, follow this rule to create complex declarations: break the complex declarations into simpler entities and use *typedef*s for those small parts and combine them together. In this case, *typedef* is very handy: One reason is that the argument as well as the return type is *sig_t*, so it looks very simple; but believe me, in reality, even with *typedef*s, some declarations can be hard to decipher!

Fortunately, there is a tool that you can use in Linux (and UNIX) when you want to understand complex declarations. It is called "cdecl" tool (c++decl for C++). You can use keywords like "explain", "declare", etc. This example is from the Linux man page for *cdecl(1)*: The input, "declare fptab as an array of pointer to function returning pointer to char" will show the output, "char *(*fptab[])()". On the other hand, you can say "explain char *(*fptab[])()", to get that English-like description. That's handy. So, what are you waiting for – just go ahead, try out some complex declarations and have fun! 

About the author:

S G Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at srganesh@gmail.com.



Up, Close,
& Personal
with the
creator of

OLPC XO

eBookReader

Internet Archive, the San Francisco-based non-profit dedicated to the universal sharing of knowledge, recently collaborated with the Massachusetts-based One Laptop Per Child Foundation (OLPC) to freely make available 1.6 million e-books to kids with the OLPC laptop (also known as the XO). The books will thus be available to around a million schoolchildren in developing countries.

To know more about the technology that makes this possible, we contacted the somewhat reclusive lead developer of the e-book reader software on the XO, **Sayamindu Dasgupta**, and got him talking about his work. Over e-mail, phone and IRC, Dasgupta shared with us details about the genesis of the software, its future and its ability to bring about change.

Q You were the i18n/l10n go-to person at OLPC; what prompted you to take a look at eBookReader?

After I became a full-time contractor for OLPC (once I had completed my graduation), my role was modified a bit—apart from i18n/l10n issues, I was also supposed to help with general software development efforts. Initially, I was looking at things like performance tuning (for example, the Browse activity being slow while drawing text, the Paint activity having redraw issues and so on), and I also took up the maintainership of a few Sugar activities.

When I was down at the OLPC office at ICC in Cambridge, I had a few informal meetings and discussions with S J Klein about book readers (we both shared a common interest), though not much happened on a concrete basis for a few months. However, in January, when Nicholas Negroponte announced a million books as a part of OLPC's refocused mission, I was asked if I could help with sprucing up the book-reader, and I jumped in.

Q Can you tell us the technical details about what you have done?

The first job was to get a decent book-reader going in our stable builds (which is based on Fedora 9). There were quite a few issues with the Read activity that shipped with 8.2.0; for example, djvu files would make the activity crash when one tried to change the zoom level. I backported a significant part of the book-reader stack (most of the heavy duty lifting is done by Evince, btw) from our experimental Fedora (rawhide) builds, and as a result the version of the Read activity in our XO OS release 8.2.1 is much more stable. (I also managed to sneak in a few features that I had been working on.)

The next phase consisted of working closely with the upstream Sugarlabs community and tying in the development process to the Sugarlabs release schedule (with

“Imagine what can happen if books have a versatile tool like Etoys, Scratch or TurtleArt built in, or imagine reading through *A Byte of Python*, and being able to try out and play around with the examples within the book itself.

backports for OLPC-specific builds and releases, whenever required). The most important feature that was added during this phase was support for EPUB, a file format that is fast becoming the de facto standard for ebook-readers worldwide (except in the case of Amazon Kindle). EPUB uses pre-existing, well-established standards to define how a book should be laid out and formatted (for example, the content itself is usually XHTML, which can be rendered via any modern browser engine).

I also started working on annotation support. This has been tricky, since we support a variety of file formats in Read activity (PDF, PS, CBZ, CBR, DJVU, EPUB, etc), but I'm trying to have basic annotation features like bookmarking, and associating notes with the bookmarks (so a user can have at least one 'side-note' per page), etc, working for all formats.

Next versions may see format-specific annotation, and perhaps, if all goes according to plan, complete editing support for certain book formats, so that one can annotate, highlight, and do whatever one wants (of course, with the option of reverting to the older, pristine version of the book if something goes wrong).

I also added some minor usability enhancements (like a nice progress bar to show how much of the book has been read while in full screen mode, a battery icon to show the remaining charge in full-screen mode, and so on).

Of course, there is a more experimental side of my work (which I did not enable in our 'production' builds). Some of it is quite obvious, like embedding video/audio inside books, etc. However, what can be really interesting is the possibility of embedding 'creative tools' within books—tools which can be then used by the children to try out new programming languages, construct their own science simulations and experiments, and so on.

People often think of Adobe Flash when they hear this—but Flash is restrictive in the sense that the consumer is limited by the design of the original demo/simulation. Imagine what can happen if books have a versatile tool like Etoys, Scratch or TurtleArt built in, or imagine reading through *A Byte of Python*, and being able to try out and play around with the examples within the book itself (an interactive Python shell embedded in a book). I have crude proof-of-concept demos of these features lying around.

Moreover, there can be interesting stuff done to a story—a seamless transition between the narrative in a book to a movie, or even a virtual world, with characters from the book coming alive around the reader (this is one of my ideas—I don't have any code to show).

Q How does this fit within the OLPC mission and vision?

I don't think I can speak on behalf of the entire OLPC community (in fact, many will have better ideas than I



The OLPC XO eBookReader

do), but here's what I think.

I firmly believe that kids can create their own content (if this sounds odd, take a look at <http://scratch.mit.edu>). However, books have been, and will be a major source of knowledge, ideas, and inspiration for generations to come (if not paper books, definitely, e-books). I tend to look back at my own childhood where books were the primary drivers of my curiosity (and perhaps the general direction my life has taken). Had it not been for that interesting book with a funny picture of a penguin on it, the one that I bought at the Calcutta Book Fair, I would probably be trying to finish my masters in Physics or English now.

When I hear of stories of people like William Kamkwamba -- a young man in Africa, with almost no education, who saw a picture of a windmill in a library book and decided to build one for his village -- my ideas about the immense potential of books to inspire are reinforced. Now imagine what could happen in a country like Uruguay, where 400,000 kids have laptops of their own, if each laptop is loaded

with a hundred books (a typical EPUB book is usually not larger than a megabyte).

Read activity (as well as Read Etexts activity—which is another great book reader by Jim Simmons) has support for sharing books. This means that if a child decides to share her book, everyone in the neighbourhood can get a copy (and not a copy that will self-destruct after 14 days). You have an effective distribution of 400,000x100 books in a country, all of them for kids!

Another interesting thing about reading books on the OLPC is the hardware. Most mainstream e-book-readers nowadays come with e-ink displays, which are extremely power-efficient, but do not support colour, and can be quite sluggish (though apparently, one can improve that). OLPC's dual mode display takes the middle path—when you go out in the sunlight to read, you turn the screen into 'reflective mode' where the screen's backlight gets turned off, and you can read in the sunlight. And when you are indoors, or need to see some colours, you just turn the backlight back on.

Q How does the software source books?

This is another part of my current efforts. Jim Simmons had originally built an Activity called Get Internet Archive Books, which queried the Internet Archive (IA) and let users download individual books into the Sugar Journal. At the same time, OLPC had been participating in the drafting process of the OPDS (Open Publication Distribution System) standard, which is an XML (more specifically, Atom) based catalogue format targeted at book-publishers and distributors. OPDS servers usually also support a standardised method for searching through the catalogues, and so I extended Jim's 'Get Internet Archive Books' activity to work with OPDS (it had to be renamed to the 'Get Books' activity, since it was no longer restricted to material from the IA).

We were one of the first to implement a client side tool for OPDS, and currently Get Books can query the Internet Archive's BookServer (which is also a very interesting project) at www.archive.org/bookserver and Feedbooks.com. (Incidentally, the Internet Archive alone has more than 1.5 million books.)

Q There aren't enough textbooks available that can make the eBookReader relevant. What do you think needs to be done?

Oh, there are quite a few interesting initiatives (and some of them are showing very promising real results as well). Some of them include CK-12 (ck12.org), Replacing Textbooks from the Earth Treasury (www.earthtreasury.org/wiki.cgi?ReplacingTextbooks), wikibooks (en.wikibooks.org/wiki/Main_Page), the work going on at OLE Nepal, etc. I'm sure to have missed out a huge number of projects—apologies to all of them. The work done by FLOSSManuals (en.flossmanuals.net) may be quite relevant as well. Local language, non-English e-text-books, are however more difficult to find.

Q How is the i18n/l10n work at OLPC shaping up?

It's in quite good shape—though, as always, we need more people who can help us out :-)

Q There is some eBookReader hardware available in the market (if not in India). Are there features that excite you?

Nook's dual display system sounds interesting—not sure how useful it would be in a real life scenario. Newspapers via e-book-readers is another interesting concept—I often dream of kids using activities to write, edit and publish their own newspaper/school journal, and circulate the final product amongst themselves by sharing (or maybe via the XS school server).

Q You aren't showing this off at events. Don't you have plans along those lines?

I plan on doing a talk at FOSS.in, Bangalore, in December. I'll also speak about the more experimental stuff in my work, as well as a general case study type of thing in a conference called 'Designing for Children' with focus on 'play + learn' at IIT Bombay in February next year.

Q You are known to be a reclusive developer. Tell us a little bit about yourself.

I am from Kolkata (West Bengal, India) where I work from home. I'm a foodie and one of my pastimes is to explore interesting food joints in my city (ranging from the roadside food-vendor to fancy restaurants). I also tend to have fun taking photographs. I have a pet cat who does not have any name (since anyone who has read Elliott's *The Naming Of Cats* is aware that it is only the cat herself who knows her own, real name). I have a habit of playing practical jokes (I was quite notorious in my college for my tricks).

Q I know that Seymour Papert, Alan Kay and a few others have been a great influence on your thoughts

“ There have been quite accurate processes for automatic conversion of music/video to digital format, but sadly the same is not true for books.

around the project. How do you conceive the 'design' part of it?


Seymour Papert's work has been significant to me in a personal manner, as it has been instrumental in my journey of trying to make sense out of my own experiences in school (which was quite a roller-coaster ride for me). Alan Kay's concept of the Dynabook has had a more direct impact on at least my work (in a manner of speaking, the entire concept of OLPC and Sugar is centred around the concepts proposed by Kay and Papert).

If you read Kay's 1972 paper *A Personal Computer for Children of All Ages*, the parallel between what is described, and what many in the OLPC and Sugar community visualise today is remarkable. The 'inspirational spark' for many of the ideas that I proposed in my 'Braindump on ebooks' (sayamindu.randomink.org/ramblings/2009/08/12/braindump-on-ebooks/) came after seeing a talk by Dr Kay in one of the meetings organised by OLPC. The concept of programmable 'microworlds', as described by Seymour Papert in his work, has also been a central theme in my ideas.

Q It is being said that e-books are the future and so forth. But with not enough content, do you think that great software like yours can make an impact? How does the relevant content get into shape?

This is an interesting question, in the sense that the transition from books to e-books involves the largest amount of 'legacy content' till date. Stored music, movies, etc, have always been, in a manner

of speaking, 'intangible' (starting from the wax cylinder to modern Blu-ray discs). But books have been firmly grounded in ink and paper, very tangible, with no need for an intermediate object or device to make the content accessible. Moreover, there have been quite accurate processes for automatic conversion of music/video to digital format, but sadly the same is not true for books (OCR software does not even exist for some of the major languages/scripts of the world).

Various efforts are being made to transform the old content into new, and some of them (Project Gutenberg's distributed proof-reading project, etc) are quite promising. The volume of pre-existing content (which cannot be reliably transformed automatically) calls for a distributed, truly worldwide effort. I have heard of a recent demo in which an application opened a book from the Internet Archive, let the reader fix errors, and then re-uploaded the changes back again. I think at some point, it may make sense to add this functionality to the reader devices and software as well.  **END**

Attribution:

- Sayamindu Dasgupta's image on the first page is copyrighted by Kushal Das. The image is released under Creative Commons Attribution-Share Alike 2.0 Generic Licence.
- The eBookReader screenshot is taken from http://sayamindu.randomink.org/profile/images/read_notes.png

By: Sankarshan Mukhopadhyay

The author has been using Fedora since the days of Fedora Core 1. He can be reached at sankarshan at fedoraproject dot org, or as sankarshan at jabber.com on IM.



Python in Research, Part 4

Optimising Returns, Minimising Risk

Can SciPy be of use in economics and social sciences? The answer, of course, is, "Yes!"



One of the best-known methods to minimise risk is insurance. However, it came into existence in the 17th century. Economic theory, like insurance, is heavily based on statistics. Economic models are not perfect, as the global recession of the last few years has shown. It is not that the modelling is wrong. Rather, more complex models are needed. The role that human behaviour plays needs to be factored in.

As models get more complex, spreadsheets with macros will not be adequate and we'll require more complex tools. An example of the possibilities is an interesting commercial product called Resolver. This blends a spreadsheet-like interface with the ability to program in Python (IronPython) and includes some support for NumPy.

It is likely that SciPy will find proponents among social scientists as well. You can be among the early converts. Explore a few trivial examples in order to get an idea of the tools available.

Reducing risk by diversification

The risk for an investment can be estimated by the standard deviation of the result. Consider the following options:

- Investment 1 at 10% return with a standard deviation of 10%

- Investment 2 at 15% return with a standard deviation of 20%

Suppose a fraction p is invested in Investment 1 and the remaining $(1-p)$ in Investment 2. What would be the resulting standard deviation? The textbook formula is:

$$\sigma^2 = p^2 \sigma_1^2 + (1-p)^2 \sigma_2^2 + p(1-p) \rho_{12} \sigma_1 \sigma_2$$

where σ_i is the standard deviation of the i^{th} investment

ρ_{12} is the correlation between the two investments

You can use matplotlib [matplotlib.sourceforge.net] to plot this equation for several values of correlation and visualise it. Figure 1 shows the plot.

```
import scipy as np
def std_dev(p, c12):
    vpsq = (p*v1)**2 + ((1-p)*v2)**2 + p*(1-p)*c12*v1*v2
    return np.sqrt(vpsq)
def avg_return(p):
    return p*r1 + (1-p)*r2
r1, v1 = 10, 10
r2, v2 = 15, 20
p = np.arange(0.1, 0.9)
# Plot the results
import matplotlib.pyplot as plt
```

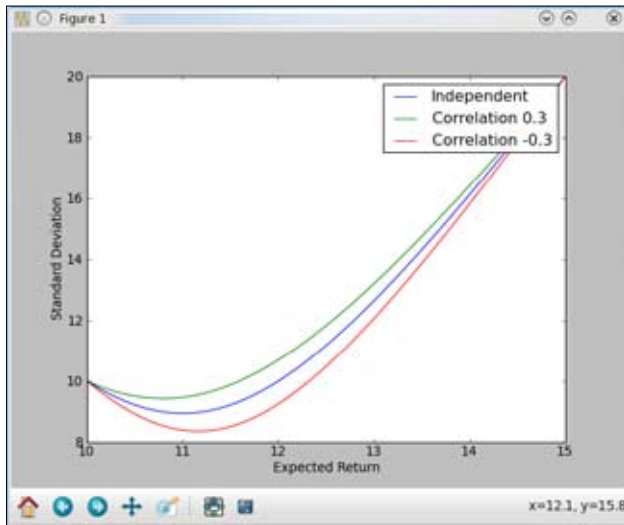



Figure 1: A graph of the standard deviation as a function of the return, based on a basket of two investments. The graphs for three values of correlation between the two investments are shown here.

```
plt.plot(avg_return(p),std_dev(p, 0),'-', label='Independent')
plt.plot(avg_return(p),std_dev(p, 0.3),'-', label='Correlation 0.3')
plt.plot(avg_return(p),std_dev(p, -0.3),'-', label='Correlation -0.3')
plt.xlabel('Expected Return')
plt.ylabel('Standard Deviation')
plt.legend()
plt.show()
```

The safest investment has a mixture of both and gives a better return than the safer of the two.

The best proportion to use will be the one that minimises the standard deviation. It is easy to get that value. Just add the following code:

```
from scipy import optimize
res = optimize.fmin(std_dev, 50, args=(0,))
print "Result", res
```

The *fmin* function in the *optimize* module accepts a function, *std_dev* in your case, as a parameter and a starting guess. Additional parameters for the *std_dev* function are passed using the *args* parameter. It then finds the optimal value. This ought to give you the result that you should have about 80 per cent in the first investment and 20 per cent in the second as the safest option, in case the investments are independent:

```
Optimization terminated successfully.
Current function value: 8.944272
Iterations: 23
Function evaluations: 46
Result [ 0.80001831]
```

You know that the investments are usually correlated. So, you may want to examine, visually, how the result varies depending upon the correlation. Add the following

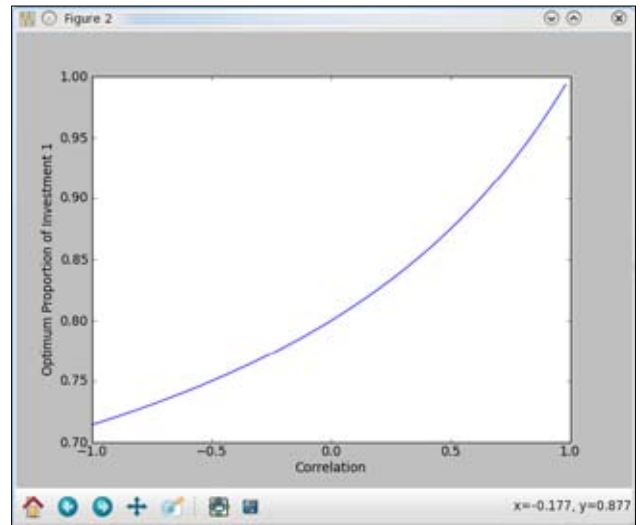


Figure 2: The graph shows how the optimum proportion for an investment varies as the correlation between two investments varies.

code to compute the optimum proportion of Investment 1 for a list of correlations and plot it:

```
def optimum_p(corr):
    res = []
    starting_val = 50
    for c12 in corr:
        res.append(optimize.fmin(std_dev, starting_val, args=(c12,)))
    # use the previous result as a starting guess
    starting_val = res[-1]
    return res
corr_values = np.arange(-1, 1., 0.02)
res = optimum_p(corr_values)
fig = plt.figure()
fig.canvas.set_window_title('Figure 2')
plt.plot(corr_values, optimum_p(corr_values), '-')
plt.xlabel('Correlation')
plt.ylabel('Optimum Proportion of Investment 1')
plt.show()
```

The result is shown in Figure 2. If both investments are perfectly correlated, then you are better off on all counts with Investment 1. However, if they have a perfect negative correlation, the best option is to invest only 70 per cent in Investment 1 and the remaining in Investment 2. An example of negative correlation is the frequently heard advice: "When stock markets are booming, interest rates fall."

If only finding the correct model in economics was easier!

LP—finding a solution

Let's suppose that the government can spend Rs 100 on increasing the consumption and production of food. It estimates that each rupee spent on rural employment will increase consumption by a rupee. This increase must be matched by an increase in production. It estimates

that each rupee spent on infrastructure will increase production by Rs 2. So, you can formulate the problem as:

```
c - p = 0 # consumption equals production
c - r = 0 # consumption equals expenditure on rural employment
p - 2*i = 0 # production equals twice the expenditure on infrastructure
r + i = 100 # sum of the expenditures must be 100
```

You can use SciPy's linalg module to solve this problem as follows, where the four columns are c, p, r, and i:

```
>>> from scipy import linalg
>>> A=[[1,-1,0,0],[1,0,-1,0],[0,1,0,-2],[0,0,1,1]]
>>> b=[0,0,0,100]
>>> linalg.solve(A,b)
array([ 66.66666667, 66.66666667, 66.66666667, 33.33333333])
```

Let's assume that the government can also waste money, w. That is:

```
r + i + w = 100
```

You now have 4 equations but 5 unknowns. So, you can use the lstsq method and get a solution:

```
>>> A=[[1,-1,0,0,0],[1,0,-1,0,0],[0,1,0,-2,0],[0,0,1,1,1]]
>>> linalg.lstsq(A,b)
(array([[ 27.27272727],
        [ 27.27272727],
        [ 13.63636364],
        [ 59.09090909]]), array([], dtype=float64), 4, array([ 2.53498934,
        1.87927105, 1.13653538, 0.86628904]))
```

The first array in the result contains the recommended values for the five parameters. The solution is undesirable, even if realistic. It suggests that 59 per cent of the money be wasted! Your model will need to minimise wastage.

Optimising in the face of constraints

SciPy, as yet, does not include a module to optimise a utility function subject to constraints. One project extending SciPy for this functionality is OpenOpt. OpenOpt uses *python-cvxopt* as one of the solvers. So, you will need to install *python-cvxopt*, which is available in the Fedora or Ubuntu repository. OpenOpt can be installed from openopt.org.

Define the optimisation problem using an *lp* example described on the OpenOpt site, as follows:

```
import scipy as np
from openopt import LP
# Parameters (c, p, r, i, w)
# Minimize w
f = np.array([0,0,0,0, 1])
# Equality constraints
# Same 4 eq as above
```


```
Aeq = np.array([[1,-1, 0, 0, 0],[1, 0, -1, 0, 0],
               [0, 1, 0, -2, 0],[0, 0, 1, 1, 1]])
beq = (0,0,0,100)
# lower bounds – all values must be >= 0
lb = [0,0,0,0,0]
# Define and solve the problem
p = LP(f, Aeq=Aeq, beq=beq, lb=lb)
r = p.solve('cvxopt_lp')
p.debug = 1
# The optimum values
print 'x_opt:', r.xf
```

If everything goes well, you should see a result that shows wastage as zero:

```
solver: cvxopt_lp problem: unnamed goal: minimum
pcost dcost gap pres dres k/t
0: 5.9091e+01 5.9091e+01 2e+02 4e-16 3e+00 1e+00
1: 3.7282e+01 3.7493e+01 3e+01 3e-14 4e-01 3e-01
2: 4.7204e-01 5.4936e-01 1e+00 1e-14 1e-02 8e-02
3: 4.7330e-03 5.5094e-03 1e-02 2e-14 1e-04 8e-04
4: 4.7331e-05 5.5096e-05 1e-04 3e-16 1e-06 8e-06
5: 4.7331e-07 5.5096e-07 1e-06 3e-14 1e-08 8e-08
6: 4.7331e-09 5.5096e-09 1e-08 3e-14 1e-10 8e-10
Optimal solution found.
istop: 1000 (optimal)
Solver: Time Elapsed = 0.1 CPU Time Elapsed = 0.01
objFunValue: 4.7331067e-09 (feasible, max constraint = 2.84217e-14)
x_opt: [ 6.66666667e+01 6.66666667e+01 6.66666667e+01
        3.33333333e+01 4.73310667e-09]
```

You can explore alternate optimisation methods—for example, maximise c, (minimise -c) or minimise (w - c), etc. Your constraints can also include inequalities well.

OpenOpt includes non-linear optimisation tools as well.

SciPy and extensions make it easy to formulate and explore solutions of the problems. Matplotlib makes it easy to visualise the results. So, you can focus on the creative and difficult part of defining the models and prevent any future crises in economics! 

References

You can learn a lot from the courses available online. For example,


- *Principles of Microeconomics*: <http://ocw.mit.edu/OcwWeb/Economics/14-01Fall-2007/CourseHome/index.htm>
- *Financial Markets*: <http://oyc.yale.edu/economics/financial-markets/>
The guest lecture on the management of the financial portfolio of Yale University is very interesting.

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com. Earlier articles in this series are available on <http://www.sethanil.com/python-for-friends>.



WHO IS GOING TO WIN THEM THIS YEAR?



Supported by:



Categories for EFY Awards*

Industrial Electronics

Resistors	Capacitors	Diodes & Transistors	Relays & Switches	Educational Kits & Trainers
EDA Tools	Solar Cells	Sensors & Actuators	Soldering & Desoldering Stations	Multimeters
Oscilloscopes	Security Systems	SMF Batteries	Inverters & UPS (Industrial)	

Consumer Electronics

Inverters & UPS (Domestic)	EPABX & Key Telephones	DVD Players	Television Sets - CRT	Television Sets - Flat Panel (LCD & Plasma)
Home Music Systems	Mobile Phones	Desktop Computers	Laptops / Notebooks	Microwave Ovens
Water Purifiers	Car Music Systems			
Electronics Man of the Year	Electronics Company of the Year	Life Time Achievement	Green Company of the Year	

Jury Awards

Sponsorship & Branding opportunities available: e-mail us: efyawards@efyindia.com

16 December, 2009 | India Habitat Center | New Delhi

For more info: www.efyawards.com

Create Cute Apps with



Part 2

This is in continuation of the article 'Create Cute Apps with Qt' that appeared in the November 2009 issue of LFY. Let's now discuss alternatives to *QSqlQuery* that are based on the model/view architecture.

In the previous article, we covered simple database connectivity concepts like connection establishment and query processing. Now, we'll cover alternatives to *QSqlQuery* that are based on the model/view architecture. However, before going into the actual details, let's understand a few more concepts about *QSqlDatabase* and *QSqlQuery* classes.

QSqlDatabase

This class is responsible for connecting to a particular database using an available driver, and provides details about the connection, driver, database, etc. Some important members of this class are:

- *addDatabase(type, name)*, which adds

a new connection with a given driver type. The optional second parameter represents the connection name, which is assumed by default as one if no name is given.

- *database(name)* returns an object for a particular connection associated with a given name. If no argument is given, a default is returned. For any further operations, processing queries and populating models, you can pass this optional object. This way, multiple connections can be maintained in a single application.
- *removeDatabase(name)* removes a particular connection.
- *close()* terminates a connection and releases any resources obtained.

- *isOpen()* checks whether a connection is active or not.
- *transaction()* begins a transaction if supported by the database and driver.
- *commit()* commits any pending operation, if any, in the present transaction.
- *rollback()* rolls back a transaction if it fails.
- *lastError()* returns an object of *QSqlError* if an error occurs.
- *connectionName()*, *connectOptions()*, *databaseName()*, *driverName()*, *hostName()*, *userName()*, *password()*, *port()*, *connectionNames()*, *drivers()* returns details of the connection, database, driver, etc.
- *setConnectOptions(opts)* sets database-specific options.
- *tables()* returns a list of tables that exist in the present database in the form of *QStringList*.
- *record(tableName)* returns the *QSqlRecord* object with field names (headers) in a specified table.

QSqlQuery

This class is responsible for executing and manipulating SQL statements. It takes a connection object as an optional argument in the constructor.

- *prepare* prepares a query with place holders. The place holder can be name-based or position based.
- *bindValue* binds the actual data to place holders.

```
query.prepare("update student set marks=? Where rollno=?");
query.addBindValue(90.23); query.addBindvalue(105);
// query.bindValue(0,90.23); query.bindValue(1,105); can be used instead
```

...OR:

```
query.prepare("update student set marks=:mks where rollno=:rno");
query.bindValue(":mks",90.23); query.bindValue(":rno",105);
```
- *exec* executes a prepared query.
- *exec(query)* executes a new query passed as a string, or you can pass this string to the constructor itself, which will be executed automatically. Once the query is executed you can use the navigation functions *first()*, *previous()*, *next()*, and *last()* to place the cursor on the first, previous, next, and last records, respectively.
- *seek(index,relative=false)* places the cursor on a record specified by the index. An absolute or relative positioning will be used based on the second argument.
- *at()* returns the present position of the cursor.
- *setForwardOnly(mode)* is to allow only forward navigation functions.
- *record()* returns the record at the current position as an object of *QSqlRecord*.
- *result()* returns the *QSqlResult* object for a given query.
- *value(index)* returns data of a field specified by the index in the form of *QVariant*.
- *size()* returns the number of rows in the result set of queries.
- *numRowsAffected()* returns the number of rows affected by the query.
- *isValid()* tells whether the cursor is on the valid record.
- *isSelect()* tells whether the current query is a SELECT

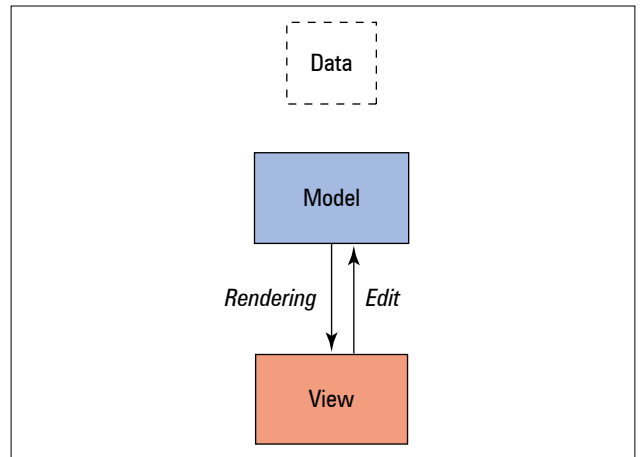


Figure 1: The model/view architecture

[Source: Qt Documentation]

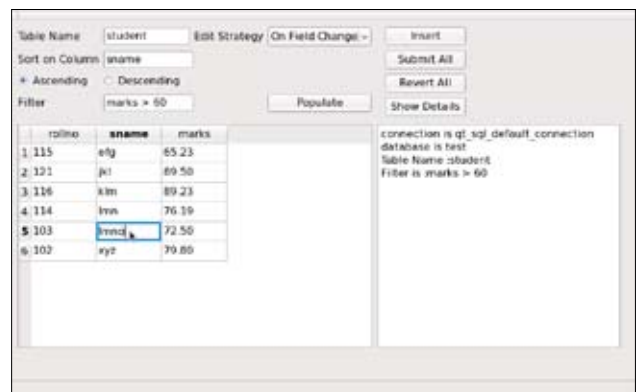


Figure 2: An editable model populated with data from the student table

operation or not.

- *lastQuery()* returns the last query executed as a string, or null, if none is executed.
- *executedQuery()* returns the last query executed successfully as a string.
- *lastError()* returns an object that represents the error state of a recent operation.

Model/view programming in Qt

The model/view architecture (see Figure 1) is a slight variation from the MVC architecture. It is achieved by combining the view and controller together. Only two layers exist here—one to communicate with the data, and the other for the user interface and interaction. Optionally, you can use delegates in between to customise the way data items are rendered and edited in a better way.

Qt comes with a set of widgets to act as models and views. Basically, three kinds of views are available, all inherited from *QAbstractItemView*. These are:

- List View
- Tree View
- Table View

Some other views like Graphics View and Web View come with specific modules. You can combine many standard models with these views. For example, you can

combine QStringList with List View, QDirModel with Tree View, etc.

The following code snippets illustrate the List View and Tree View, assuming two widgets by the object names lv and tv are added to the UI form:

```
QStringList qsl = db.tables(); // db is connection object returned after
establishment
QStringListModel *qsm1=new QStringListModel();
qsm1->setStringList(qsl);
ui->lv->setModel(qsm1);

QDirModel* qdm=new QDirModel();
ui->tv->setModel(qdm);
ui->tv->setRootIndex(qdm->index("/home"));
```

The QSql module comes with three types of models—QSqlQueryModel, QSqlTableModel and QSqlRelationalTableModel. These SQL models are combined generally with the Table View. In this article we'll discuss QSqlQueryModel and QSqlTableModel.

QSqlTableModel

This is derived from QSqlQueryModel. It provides a model to display and edit the data that belongs to a single table in a particular database. We'll explain the creation of this model and some of its important members using the following example.

Create a Qt GUI project with the added module QSql, and design the interface as shown in Figure 2. A table with the name student and fields *rollno*, *sname* and *marks* is assumed in the *test* database of MySQL.

We'll use the following names for the various widgets:

- Labels – *lblTableName*, *lblSortOn*, *lblFilter*, *lblEditStrategy*
- Line Edits – *leTableName*, *leSortOn*, *leFilter*
- Radio Buttons – *rbAscending*, *rbDescending*
- Combo Box – *cbEditStrategy*
- Push Buttons – *pbPopulate*, *pbInsert*, *pbSubmitAll*, *pbRevertAll*, *pbShowDetails*
- Table View – *tvData*
- Text Edit – *teDetails*

Add the following members to the header file:

```
QSqlDatabase db;
QSqlTableModel* model;
```

Fill the constructor with the code as discussed in the previous article for connection establishment.

Here's the code for the *Populate* button:

```
1 model = new QSqlTableModel;
2 model->setTable(ui->leTableName->text());
3 int sortIndex=model->fieldIndex(ui->leSortOn->text());
```

```
4 if(ui->rbAscending->isChecked())
5     model->setSort(sortIndex,Qt::AscendingOrder);
6 else
7     model->setSort(sortIndex,Qt::DescendingOrder);
8 model->setFilter(ui->leFilter->text());
9 switch(ui->cbEditStrategy->currentIndex())
10 {
11     case 0:model->setEditStrategy(QSqlTableModel::OnFieldChange);
12     break;
13     case 1:model->setEditStrategy(QSqlTableModel::OnRowChange);
14     break;
15     case 2:model->setEditStrategy(QSqlTableModel::OnManualSubmit);
16     break;
17 }
18 model->select();
19 ui->tvData->setModel(model);
```

Here is the description for the above code:

- In Line 2, the *setTable* method links the model with the specified table. It takes the connection object as an optional second argument.
- In Line 3, *fieldIndex* returns an index for the given field name in the form of an integer.
- In Lines 5 and 7 the *setSort* method sorts the table data on a given column with a given order. (You should call *setSort* before population. The sort function can only be used when the model is populated with table data.)
- In line 8 *setFilter* takes the where clause of SQL query as string and filter records.
- In lines 11, 13 and 15 we're setting the edit strategy as *onFieldChange*, *onRowChange* or *Manual*.
- In line 18 the *select* function populates the model with data.
- In line 19 *setModel* combines model and view.

Code for *Insert*:

```
model->insertRows(model->rowCount(),1);
```

After the last row, one empty record will be added. Enter some data in the empty cells and submit the operation.

Code for *Submit All*:

```
model->submitAll()
```

This operation commits all pending operations, if the edit strategy uses the manual mode. We can use the *submit()* function instead to commit any changes on a per-field or per-row basis.

Code for *Revert All*:

```
model->revertAll()
```

This rolls back all changes. Instead, you can also use the *revertRow(rowNum)* to discard any changes in the specified row and *revert()* to rollback any changes in the current row.

Code for *Show Details*:

```
ui->teDetails->clear();
ui->teDetails->append("Connection is : "+
    model->database().connectionName());
ui->teDetails->append("Database is : "+
    model->database().databaseName());
ui->teDetails->append("Table Name : "+model->tableName());
ui->teDetails->append("Where Clause : "+model->filter());
```

QSqlQueryModel

This is a base class for QSqlTableModel and provides read-only access to the table data:

```
QSqlQueryModel* model = new QSqlQueryModel;
model->setQuery("select * from student");
model->setHeaderData(1,Qt::Horizontal,QVariant("Name"));
```

We use the *setQuery* function instead of *setTable* here. Unlike the table model, we don't need to call the select function explicitly. The data is automatically populated using the *setQuery* call. The *setQuery* function takes the connection object as an optional second argument. We also use the *setHeaderData* function to rename column headers. In this example, *sname* is renamed Name.

We can bind the above example to the table view as well, in a similar fashion.

Other applicable operations for QSqlQueryModel are:

- *clear()* to detach from the table and release all acquired information.
- *removeColumns(fromCol, numCols)* filters a specified number of columns from a given column.
- *record(pos)* returns the QSqlRecord for record at a given position.
- *setHeaderData(index,orientation,value)* renames the row, column headers as per given orientation which is horizontal per columns and vertical per rows.
- *headerData(index,orientation)* fetches the row or column header names.
- *lastError()* returns an error object if there is any failure in the recent operation.

We can use these models without mapping to View.

For example, here's the code to get all data associated with these models:

```
int rollno;QString name;double marks;
for(int i=0;i<model->rowCount();i++)
{
    rollno=model->record(i).value(0).toInt();
    name=model->record(i).value(1).toString();
    marks=model->record(i).value(2).toDouble();
    qDebug() << rollno << "," << name << "," << marks;
}
```

And here's the code to replace or insert a particular record using QSqlTableModel:

```
QSqlRecord rec;
```

```
QSqlField fd1("rollno",QVariant::Int); fd1.setValue("121");
rec.insert(0,fd1);
QSqlField fd2("sname",QVariant::String); fd2.setValue("jkl");
rec.insert(1,fd2);
QSqlField fd3("marks",QVariant::Double); fd3.setValue("89.5");
rec.insert(2,fd3);
model->insertRecord(-1,rec);
model->setRecord(5,rec); // to replace 5th record with above one
```


QSqlError

To ensure any database operation has not failed, the *lastError* function should be called, which returns an object of the QSqlError class. Then, checking the *isValid()* function on that object will indicate whether an error has occurred or not. Some important members of the QSqlError class are:

- *type()* returns the exact type of error—connection error, statement error, transaction error, unknown error, or no error.
- *number()* returns database-specific error numbers.
- *databaseText()* returns the error message as reported by the database.
- *driverText()* returns the error message as reported by the driver.
- *text()* returns the concatenated string of the database and driver error messages.

The following code depicts the usage of the QSqlError function:

```
//some operation with a model or query object
QSqlError err=model.lastError();
if(err.isValid())
{
    qDebug() << err.databaseText();
    qDebug() << err.driverText();
    qDebug() << "Error Type" << err.type();
    qDebug() << "Error no: " << err.number();
    // Do error handling
}
```

In this article, we haven't listed or discussed the details of some members. Referring to the QtSql section of the Qt documentation can help you further. 

References

- Official Qt website: qt.nokia.com
- Documentation for Qt and related products: doc.trolltech.com

By Rajesh Sola

The author is a faculty member of the Computer Science Department at NBKRIST, Vidyanagar. He is a contributor to the OpenOffice.org project and is keen on promoting FOSS awareness and adoption in rural areas. He believes in encouraging and supporting students to take the open source road. You can reach him at rajesh at lisor dot org.



Sandya Mannarswamy

This month's column contains the solutions to some of the programming-related interview questions we featured in the last issue.

Many readers have written back with answers to the questions we featured in last month's column. Others have mentioned that they found some of the questions challenging. So we will keep some of the difficult questions open till next month, giving readers more time to come up with the solutions themselves. However, in this month's column, let us discuss some of the easier ones. Please do send your answers to the other questions so that I can feature your solutions in next month's column.

(1) Can you have both const and volatile qualifiers applied to a single declaration in C? If so, can you give an example? If not, why not?

The surprising answer is that both const and volatile can be applied to a single declaration. Const and volatile are type qualifiers. Const means that the program cannot modify the value of a given variable qualified by a const qualifier. For instance, `const int NUM = 100` is a valid C declaration, which means that NUM is an integer variable whose value cannot be modified by the program. The qualifier volatile means that the value of the variable can change by means that are external to the program, such as by means of an I/O device.

Consider a variable declared as const volatile int hw_register; The const declaration means that the value of hw_register can not be changed by this program, whereas the 'volatile' specifier implies that its value can be changed at any time from outside the program. However, remember that 'volatile' prevents the compiler from many of its optimisations and hence an injudicious use of 'volatile' can have an impact on your application's performance.

(2) Consider the following code snippet:

```
main()
{
    int a, b;
    a = foo();
    b = foo();
    printf("a = %d b = %d\n", a, b);
}
```

You do not know anything about function 'foo' except that it returns an integer. Are the values of *a* and *b* that are printed by the `printf` function, the same? If not, can you give an example of a function *foo* that will cause *a* and *b* to be different?

No, you can't make any assumption about the values

returned by *foo* by the two invocations. For instance, consider the following definition of *foo*():

```
int foo(void)
{
    static int counter = 0;
    return counter++;
}
```

In this case, the first invocation of *foo* will return 0 and the second invocation will return 1. Since you do not know how the internal state of *foo* will be affected by each of its invocations, we cannot make any assumptions about the values returned by the two successive invocations of *foo* with the same value of arguments, though they come immediately, one after the other, with no intervening statements.

(3) Given a word, can you find all the anagrams of that word? If you were asked to solve this problem for only one word, what would your solution be? If you were asked to find the anagrams for 10,000 words, would your solution change?

This is a problem we have discussed in one of the earlier columns. The brute force solution of finding the number of anagrams consists of forming each permutation of the word and then looking up that particular permutation in a dictionary to see if it is a valid word. While this may be possible for a single word, repeating this for thousands of words will take too long to be practical. Hence the best thing would be to first build a set of anagram classes of your dictionary so that given a word, you can look up the dictionary for all its anagram classes in a single shot.

This problem is discussed in detail in Chapter 2 of '*Programming Pearls*' by John Bentley. The basic idea is to create a signature for each word by sorting it in alphabetical order. For example, both the words 'pans' and 'span' have the signature 'ansp'. Then sort all the words by their signatures to form the anagram classes.

(4) What is the time complexity of searching for an element in each of the following:

- (a) A linked list containing *N* elements
- (b) A hash table containing *N* elements
- (c) A binary search tree containing *N* elements
- (d) A binary heap containing *N* elements
- (e) A d-heap containing *N* elements

In a linked list, if you are looking for an element that is not

present in the list, you have to walk through all elements in the list before you can say for sure that the element is not present in the list. Hence, the complexity of search on a linked list is $O(N)$. For a hash table, assuming no collisions, the complexity is $O(1)$. For a binary search tree (which we assume to be balanced), it is $O(\log N)$. For a binary heap and a d-ary heap, it is $O(N)$. I leave it to the reader to figure out why it is $O(N)$ for a heap.

(5) Consider the following code snippet:

```
int find_fib (int N)
{
    assert (N > 0);
    int f1 = 1;
    int f2 = 1;

    if (N < 2)
        return 1;

    return (find_fib(N-1) + find_fib(N-2));
}
```

Given that *find_fib* is called from the main with N as 25, how many total calls to *find_fib* occur?

A call to *find_fib(N)* ends up calling *find_fib(N-1)* and *find_fib(N-2)*. Note that a call to *find_fib(N-1)* will end up calling *find_fib(N-2)*. Hence we find that *find_fib(N-1)* gets called once, *find_fib(N-2)* gets called two times, *find_fib(N-3)* gets called once in *find_fib(N-1)*, once each by every one of the two calls to *find_fib(N-2)* and hence *find_fib(N-3)* gets called three times in total. We can write a recurrence relation to the number of calls to *find_fib* as:

$$C(N) = C(N-1) + C(N-2) + 1$$

...where $C(N)$ is the number of calls to *find_fib* and $C(0) = C(1) = 1$. I leave it to the reader to solve the recurrence and work out the value for $N = 25$.

Sorting/searching/string manipulation problems

(6) You are given an array of $2n+1$ integers and told that except for only one element, all other elements have a duplicate in the array. Can you find the one element that has no duplicate in the array? What is the time complexity of your solution?

While a brute force solution would be to sort the array and perform a linear scan to find the duplicate, there is a much quicker solution. Remember that exclusive ORing of an element 'A' with itself is 0. Since all elements except one have a matching duplicate, just do a XOR of all elements. The unmatched element will be the result.

(7) You are given an array A of N integers. A majority element M appears more than $N/2$ times in the array. For instance, given the array of integers, 10, 6, 10, the majority element is 3, 10, 10. Given the array 10, 6, 10, 3, 10, 21, 5, there is no majority element. For the given array of N integers, write an algorithm to find whether a majority element exists and if so, what it is. The algorithm should have a worst-case complexity of $O(N)$.

While at first glance it looks like sorting the array and doing a linear scan to find the majority element would be the right approach, this solution will have a complexity of $O(N \log N)$, which is larger than the $O(N)$ required. Hence we need to devise a more clever approach.


There is a well-known algorithm that achieves the required $O(N)$ complexity by approaching the problem in two steps. The first step is to call the function *find_possible_majority_element*, which returns an array element that could possibly be the majority element in the given array A. The algorithm *find_possible_majority_element* returns an element from the given array A. In the second step, the algorithm *CheckIfMajorityElement(int elem)* actually finds out whether the element *elem* returned by *find_possible_majority_element* in the first step is indeed the majority in A by counting the number of occurrences of *elem* in array A.

The function *find_possible_majority_element* takes the first element of the array A and puts a score 1, and considers it as the current *majority* element, *elem*. It then takes this element, *elem*, and the current score of 1 as the basis and traverses through the other elements of the array, performing the following checks:

- If the next element in the array is equal to *elem*, then the score is incremented by 1.
- If the next element in the array is not equal to *elem*, then the score is decremented by 1.
- If at the next element, the current score is equal to zero, then the score is set to 1 and this element is considered as *elem*.
- After traversing all elements, finally *find_possible_majority_element* returns *elem* as the possible majority element.

In the next step, *CheckIfMajorityElement* takes this element *elem* as the input argument and counts the number of occurrences of the element in the array A. If the count is greater than $n/2$, this element is the majority element; otherwise, null is returned. I leave this as an exercise for readers to write the actual code for the solution.

We will discuss the solutions to the remaining problems in our next column. Meanwhile, I would like to request readers to send in their solutions to these questions. Also, if you are preparing for computer science programming interviews, you may find it useful to visit my programming interviews discussion group "*Computer Science Interview Training (India)*" on LinkedIn.

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me along with your solutions and feedback at sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming!  **END**

About the author:

Sandya Mannarswamy. The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.

SubConf is the annual conference of the Subversion community. SubConf 2009 was the third such event, held at Munich, Germany, from October 27 to 29, 2009. While it's a user conference where Subversion users from various parts of the world participate, it does feature developer hackathons where the project's core developers come together to discuss the roadmap, hack on code, et al. Developers also meet the users to get feedback, and study their requirements so that future releases can cater to these.

SubConf 2009 had 10 core Subversion developers at the conference venue—Stephen Butler, Stefan Sperling and Neels Hofmeyr (of Elego); Julian Foad of WANDisco; Greg Stein (a popular open source developer); Hyrum K. Wright (Subversion Corp), Lieven Govaerts, Bert Huijben (of The Competence Group), C. Michael Pilato and myself (from Collabnet, Inc).

All of us were locked up in a room (hackathon) at the conference hotel for three days, where we discussed various things related to Subversion development such as the Working Copy Next Generation (WC-NG) library, usage of a scratch pool, itpool in Subversion code base, the release roadmap, interesting issues to work on, etc. Of course, the hackathon was not just a discussion forum; we also undertook some really productive programming—there were approximately 70 commits to the Subversion

repository with close to +46696/-36666 lines of change!

The first day of the conference started officially in the evening around 7:00 pm with the Subversion Round Table, where users from various organisations posted their queries and feedback about the version control system (VCS). They also explored the possibilities of getting some features introduced in future releases of the VCS. This was a fruitful discussion. Every year this results in bringing in new requirements to the project, directly from the target audience.

On the second day of the conference we had many talks regarding VCS. Long-time Subversion developer (from Jan 2001) C. Michael Pilato delivered the keynote address. He spoke on the history of Subversion, the way the community works, why CollabNet chose to make Subversion an open source project, etc. The talk covered the legacy of the Subversion Community and the advancements it had gone through over the years.

Some of the talks delivered on the second and third day of the conference were as follows (there were even more talks, but they were not in English. See reference [4], [5] for presentation slides link):

- *Subversion Release Process* by Subversion release manager Hyrum Wright and Stefan Sperling
- *Bringing Subversion to the Java World* by Alexander Kitaev and Alexander Sinyushkin
- *WC-NG* (Subversion's new working copy management library) by Hyrum Wright
- *Comparing Apples to Oranges—Subversion, git and Mercurial* by Stefan Sperling and Stephen Butler
- *Moving from SVN to Mercurial* by Zsolt Koppany and Janos Koppany

SubConf 2009

A Report

Senthil Kumaran S, a full committer to the Subversion project, shares his experiences at the three-day conference held in Munich recently.

- *Server Side Java Bindings for Subversion* by Dave Brown
- *SVN Obliterate* by Julian Foad
- *Coding Control* by Perforce Software's Tony Smith


The message from Subversion developers to the users was to do some real testing of the pre-release versions (we don't want you to try on production data, though) of Subversion software to catch bugs early. The fact is it's difficult for developers to try and mimic the varied environments in which Subversion is deployed in organisations, and only the users can chip in here.

Users also expressed a requirement for pre-release binaries other than the source tarballs (which, currently, the Subversion Community is not engaged in providing). The developers took note of this and expressed a desire to work on some mechanism to get this done in the future. For what it's worth, in the recent past, the project has already started providing nightly tarballs of the latest trunk development sources (see reference [6]).

Another interesting take away from the conference was the community's feeling about Distributed Version Control Systems (DVCS). The community is excited about DVCS, since we are part of advancing the 'state of the art' and we are happy that, ultimately, we have competitors in the version control world. With the latest improvements on the WC-NG library, Subversion will be able to get features like offline commits, shelving, etc. I'd say that it's

premature to talk about this aspect now, but it is possible in the foreseeable future.

It was a nice experience for me to lurk around with the other developers at SubConf, with whom I've only communicated by e-mail till now. We also had a surprise in the week after the conference with the announcement made at ApacheCon 2009, about the Subversion project finding a new home at the Apache Software Foundation!

Subversion is a widely used version control system in many free software projects and corporate environments. The latest release of Subversion (1.6.6) source can be downloaded from <http://subversion.tigris.org>. If you want Subversion binaries for different platforms visit <http://open.collab.net>. 

References

- <http://www.linuxforu.com/news/subconf-2009-a-report-day-1>
- <http://www.linuxforu.com/news/subconf-2009-a-report-day-2>
- <http://www.linuxforu.com/news/subconf-2009-a-report-day-3>
- <http://2009.subconf.de/vortraege/1-konferenztag>
- <http://2009.subconf.de/vortraege/2-konferenztag>
- <http://orac.ece.utexas.edu/pub/svn/nightly>
- <http://www.subconf.de>
- <http://www.collab.net>

By: Senthil Kumaran S.

The author is currently employed by CollabNet, and is working for the Version Control group. He is a 'full committer' of the Subversion project and a free software enthusiast. To know more, visit www.styleesen.org.



▲ HOOD (conference organizers) stall



▲ CollabNet stall at SubConf 2009



▲ Perforce Stall at SubConf2009



▲ Subversion Hackathon Room at SubConf 2009



Version Compatibility

Here's some advice on choosing the right release of Subversion.

The toughest job for any server administrator is to choose the correct version of software to install and use for maintenance. Most server software have corresponding clients that access the software to get data from them. This kind of client-server model creates a few problems when a server administrator is trying to choose the right version of server software.

In this article, we will discuss one such problem in choosing the appropriate release of the popular version control system, Subversion. This article will also help decipher version compatibility among most of the free software available.

Servers

All the Subversion servers that have the same major version number will be compatible with each other. For example, if you are using Subversion 1.6.3 and want to switch to 1.6.x, which is greater than 1.6.3, both are compatible without any modifications. On the other hand, when you want to switch from 1.5.x to 1.6.x server, there may be some difference in the repository *backend/fsfs/bdb*, for which

the Subversion community provides tools to make this upgrade smooth without any problems. The Subversion servers that have a major version number will be compatible with all Subversion clients, which will have the same major version number.

Clients

When the minor version number (which is '6' in version 1.6.4) of a Subversion client changes, you'll typically notice more features added to the Subversion client, but without breaking the backward compatibility between the minor version numbers. In other words, none of the features available in the same minor version number series will be broken when we go for an upgrade. All the features in 1.5.x will continue to be available in 1.6.x along with some additional features.

The Subversion clients and servers will have sub minor version numbers such as the 4 in version 1.6.4. These sub minor versions are released with bug fixes. There is a good chance that these bug fixes are also back-ported to earlier releases, such as 1.5.x or 1.4.x, if these branches are still being maintained. But none of the features

introduced are backported to a previous minor version release. For example, tree conflicts introduced in 1.6.0 are not backported to the 1.5.0 release branch.

Let us take the case of the Subversion 1.6.x release and see how it gets affected between the bug fix releases. When Subversion 1.6 was officially released, it got the version number 1.6.0 with a number of new features, such as the following:

- Repository root relative URLs
- Detection of tree conflicts
- cTypes Python bindings
- Logging support for *svnserve*
- Support for file externals
- Sparse directory exclusion

Apart from the features introduced in this 1.6.0 release, there were many bug fixes and improvements to the existing features. Then, once in two months, Subversion developers started releasing 1.6.1, 1.6.2, etc. These included improvements and bug fixes on existing features and the features introduced in the 1.6.0 release.

There was a special case in the 1.6.x release branch, where the Subversion developers came up with a release 1.6.4, which was for just one security fix. In this release, bugs weren't fixed, but there was a single patch to fix a security issue in Subversion clients/servers. This could have been fixed as a part of a bug fix release in Subversion 1.6.4 but the developers did not want to wait for all the bug fixes scheduled for 1.6.4 to get in before they released the patch, which would also have to undergo a soak period of two months. The complete bug fix release that subsequently came out was version 1.6.5.

In any release, if you upgrade to the same minor version—that is, 6 in 1.6.4 or 1.6.5—you will not experience any difference in the software. The dependencies that worked for 1.6.0 will continue to work for 1.6.4, 1.6.5, etc, unless you decide to change the dependencies to a different version due to some


security reasons.


Hence, a user can choose the following paths to upgrade to 1.6.6, which is the latest version, as of writing this article:

- 1.6.0 to 1.6.6
- 1.6.1 to 1.6.6
- 1.6.2 to 1.6.6
- 1.6.3 to 1.6.6
- 1.6.4 to 1.6.6
- 1.6.5 to 1.6.6

This also applies to changing from any arbitrary version to one with the same minor number, like 1.6.2 to 1.6.4 and various similar combinations.

In case of Subversion 1.x.x and 2.x.x, all the above-mentioned compatibilities may or may not work. The Subversion community does not promise to preserve the backward compatibility between major versions of Subversion, in which case a 1.x.x client will not work smoothly with a 2.x.x server. Similarly, a 1.x.x server will not work smoothly with 2.x.x clients. Though we do not have a 2.x.x version of a Subversion release yet, this was the case with the 0.x.x Subversion release and the 1.x.x Subversion release.

 **Note:** Subversion 1.6.4 and 1.5.7 are important security fix releases and all servers must be updated for this fix, if Subversion servers using HTTP, HTTPS or the SVN protocols are exposed to the public.

The latest release of the Subversion (1.6.6) source can be downloaded from <http://subversion.tigris.org>. If you want Subversion binaries for different platforms, visit <http://open.collab.net>. 

By: Senthil Kumaran S.

The author is currently employed by CollabNet, and is working for the Version Control group. He is a 'full committer' of the Subversion project and a free software enthusiast. To know more, visit www.styleesen.org.

Your favourite Linux Magazine
is now on the Web, too.
LinuxForU.com

Follow us on Twitter @LinuxForU

THE COMPLETE MAGAZINE ON OPEN SOURCE

LINUX
BETA
ForYou



A Voyage to the Kernel



Part 19

Day 18: Segment 3.8

*A*s I promised in the last issue, in this article, we will explore networking.

OSI reference model

The OSI (Open System Interconnection) reference model is a multi-layered computer network protocol architecture that has seven distinct layers from top to bottom (refer to Figure 1). These are the application, presentation, session, transport, network, data and physical layers. After covering some of the unique features of the architecture, we will move on to networking in Linux systems.

The whole purpose of this division is to subtly divide the process of networking. Conceptually, the layer below a particular layer is therefore one layer to provide support to the former. By definition, you can have a 'link' between two instances belonging to a particular type of layer. You can assume that there is a kind of *horizontal protocol* connection existing between them.

Let's discuss these layers from the bottom to the top. The physical layer, as the name suggests, is what deals with the hardware device and the physical medium. This obviously covers items like hubs, repeaters, network adapters, etc. You can also place *cable specifications* and *pins* under this group.

The data layer deals mainly with the *procedural means* to transfer data between the network and the other elements in the network. This may perplex some people. Well, in order to avoid that confusion, you can assume that the physical layer simply meddles with the linking of a single entity with the physical medium, while the layer above it handles multiple devices. The data layer has the ability to perform error correction and control the flow of data.

The network layer deals with the transferring

of data sequences of different lengths from a given source to the intended destination. This layer does the routing functions and is also responsible for fragmentation and reassembly of data in many cases. The network layer can report delivery errors as well; and these facilitate the operation of routers at this layer. The Internet Protocol (IP) is the most well known among those belonging to this layer.

The transport layer is responsible for maintaining the quality of service. It does so by requesting this quality of the layer below it and, in turn, helping in reliable data transfer. The layer can also perform segmentation and de-segmentation processes. Just like the data layer, this layer plays a role in flow and error control as well. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are two examples that belong to this category.

The session layer, obviously, deals with the handling of sessions. It manages the connection between two computers (say, a local and remote application). The presentation layer is actually a link layer that does some sort of *translation* service. It takes inputs from the higher layer and feeds them into the session layer in such a way that the layer can handle them.

The top-most layer is the application layer that interacts with software applications and is architecturally positioned close to the end user. Typical examples could be Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), etc.

In short, you can see that in the OSI model, the N^{th} -layer is supported by an $(N-1)^{\text{th}}$ layer, which helps the former layer to enable error-free transfer of data. These architectural changes were introduced as per the demand.

The Internet has expanded a lot over the last

15 years. If you look at the ARPANET logical map (see Figure 2), you can see how simple (!) the network was. Back then, nobody thought that we might run out of IPv4 unique values and have to go for alternate options like IPv6. This indicates that the layers can change in future to improve the performance of the network.

Though I mentioned TCP and UDP while discussing Layer 4, I think it is necessary to discuss them a bit more. The TCP/IP suit is one of the core protocols we use in the Internet. The IP is largely responsible for the delivery of the data by taking it all the way through the network. But TCP is concerned with issues only at the source and at the destination. A typical example could be a Web server and a browser. Please refer to Figure 3 for information concerning its header.

UDP can be employed to send messages to other hosts on the network without the requirement that *earlier communications should set up special transmission channels*. Its header format is shown in Figure 4.

One thing that you may notice is that since UDP doesn't require a *formal setting up process* for transmission, it affects the reliability and integrity of the data transmitted. If you wonder why we need that, then I would suggest you think about services that are real-time and those applications in which the waiting for lost packages is not the preferred option.

If you wish to explore more about networking in general, I would suggest you read *Computer Networks* by Andrew Tanenbum.

Linux networking

Let's ponder a little more about networking in the Linux platform. First of all, we will look at some basic commands.

ifconfig: This is used to configure the *kernel-resident network interfaces*. If you wish to display the status of all interfaces (including those that are down), you can issue the following command (please don't skip the results, as that may help you get accustomed to the available interfaces):

```
aasisvinayak@GNU-BOX:~$ ifconfig -a
eth0  Link encap:Ethernet HWaddr 00:25:64:56:f4:e1
      inet addr:---HIDDEN--- Bcast:131.227.156.255 Mask:255.255.255.0
      inet6 addr: ---HIDDEN--- Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:371084 errors:0 dropped:0 overruns:0 frame:0
      TX packets:245076 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:476799041 (476.7 MB) TX bytes:22141513 (22.1 MB)
      Interrupt:29 Base address:0x6000

eth2  Link encap:Ethernet HWaddr 00:26:5e:7d:8d:53
      inet6 addr: fe80::226:5eff:fe7d:8d53/64 Scope:Link
      UP BROADCAST MULTICAST MTU:1500 Metric:1
      RX packets:0 errors:0 dropped:0 overruns:0 frame:0
      TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
      Interrupt:17 Base address:0xc000

lo    Link encap:Local Loopback
```

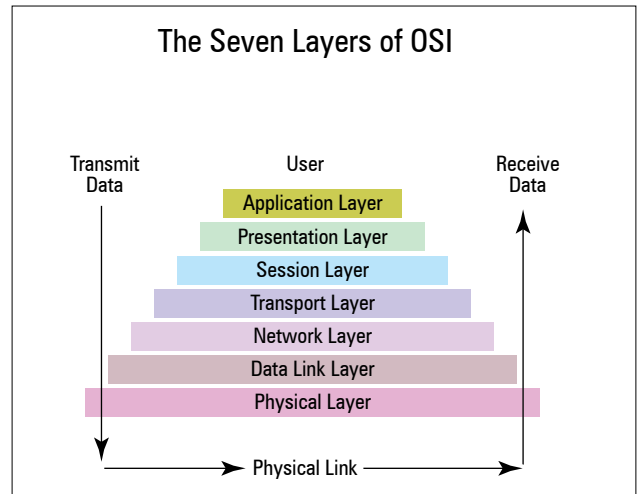


Figure 1: The seven layers of OSI

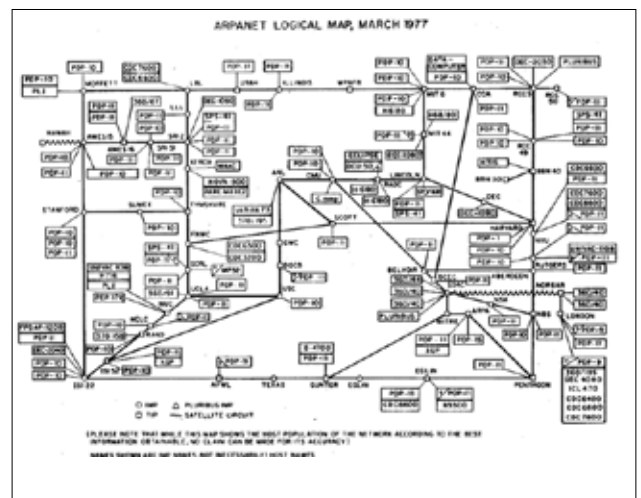


Figure 2: ARPANET logical map, March 1977

(Source: The Computer History Museum)

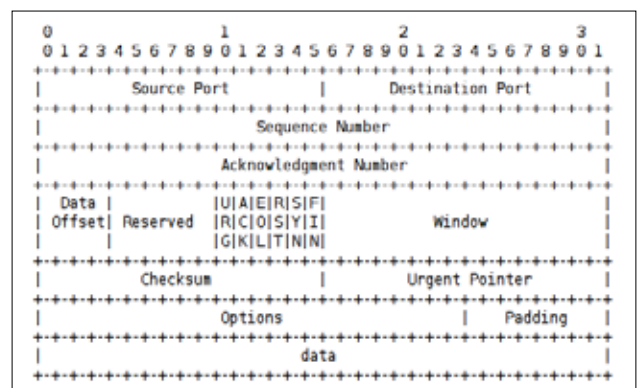


Figure 3: The TCP/IP header format

(Source: RFC761)

```
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
Interrupt:17 Base address:0xc000
```

```
lo    Link encap:Local Loopback
```

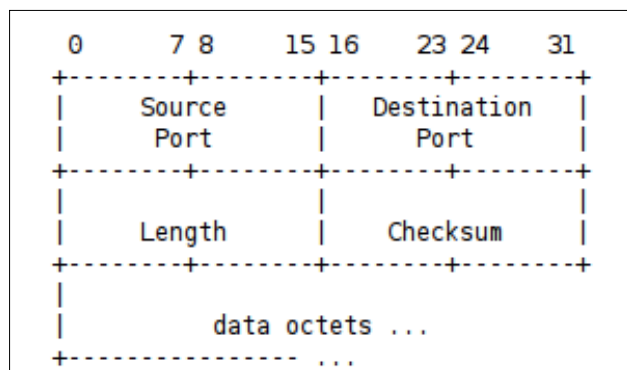


Figure 4: The UDP header

(Source: RFC768)

```
aasisvinayak@GNU-BOX:~$ cat /proc/net/route
Iface Destination Gateway Flags RefCnt Use Metric Mask M
TU Window IRTT
eth0 000CE3B3 00000000 0001 0 0 1 00FFFFFF
0 0 0
eth0 0000FEA9 00000000 0001 0 0 1000 0000FFFF
0 0 0
eth0 00000000 019CE3B3 0003 0 0 0 00000000
0 0 0
aasisvinayak@GNU-BOX:~$
```

Figure 5: The `/proc/net/route` file

```
aasisvinayak@GNU-BOX:~$ netstat -r
Kernel IP routing table
Destination Gateway Genmask Flags MSS Window irtt Iface
131.227.156.0 * 255.255.255.0 U 0 0 0 eth0
link-local * 255.255.0.0 U 0 0 0 eth0
default resnet-156.surr 0.0.0.0 U0 0 0 eth0
aasisvinayak@GNU-BOX:~$
```

Figure 6: The output for `netstat -r`

```
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:974 errors:0 dropped:0 overruns:0 frame:0
TX packets:974 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:153795 (153.7 KB) TX bytes:153795 (153.7 KB)

pan0 Link encap:Ethernet HWaddr 66:a2:20:ba:f3:10
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)

vboxnet0 Link encap:Ethernet HWaddr 0a:00:27:00:00:00
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
```

You may use the `down` option to shut down an interface, and `arp` to enable (or disable) the use of the ARP protocol on an interface. `io_addr addr` is another option to start an address in the I/O space. And if you

wish to assign an IP address, you may just employ the `address` option.

Displaying the output (content) of `/proc/interrupts` and analysing it is quite handy in some cases (say, if you want to look at the number of interrupts per IRQ). You can do this by issuing the following commands:

```
aasisvinayak@GNU-BOX:~$ sudo cat /proc/interrupts
[sudo] password for aasisvinayak:
CPU0 CPU1
0: 2275331 1552933 IO-APIC-edge timer
1: 917 457 IO-APIC-edge i8042
8: 0 1 IO-APIC-edge rtc0
9: 9 4 IO-APIC-fastEOI acpi
12: 79030 76412 IO-APIC-edge i8042
17: 4193 2697 IO-APIC-fastEOI eth2
20: 457 159 IO-APIC-fastEOI ehci_hcd:usb2, uhci_hcd:usb3,
uhci_hcd:usb6
21: 83736 41592 IO-APIC-fastEOI uhci_hcd:usb4, uhci_hcd:usb7,
HDA Intel
22: 27 13 IO-APIC-fastEOI ehci_hcd:usb1, uhci_hcd:usb5,
uhci_hcd:usb8
28: 79298 55301 PCI-MSI-edge ahci
29: 299499 306516 PCI-MSI-edge eth0
30: 181527 188303 PCI-MSI-edge i915@pci:0000:00:02.0
NMI: 0 0 Non-maskable interrupts
LOC: 1551082 2026604 Local timer interrupts
SPU: 0 0 Spurious interrupts
CNT: 0 0 Performance counter interrupts
PND: 0 0 Performance pending work
RES: 789579 767812 Rescheduling interrupts
CAL: 156 160 Function call interrupts
TLB: 5216 7330 TLB shutdowns
TRM: 0 0 Thermal event interrupts
THR: 0 0 Threshold APIC interrupts
MCE: 0 0 Machine check exceptions
MCP: 34 34 Machine check polls
ERR: 0
MIS:
```

If you wish to look at the static routing table, you can `cat` the `/proc/net/route` file. Figure 5 shows the output on my PC.

You might have tried the `netstat` command as well. But the above commands will work even if you don't have the `netstat` utility. (I have seen some changes in 2.4 and 2.6, but you need not worry about this now.)

Since most of the *distros* have this utility, you can issue `netstat -r` (or `-nr`) to obtain the routing table—readability is better in this case (refer to Figure 6). You can also use `route -n` for the same purpose.

`/etc/services` is a very vital file that helps you find how the port numbers are linked to the named services. For your reference, here is a standard entry in the file:

```
tcpmux 1/tcp # TCP port service multiplexer
```



```

echo      7/tcp
echo      7/udp
discard   9/tcp      sink null
discard   9/udp      sink null
systat    11/tcp     users
daytime   13/tcp
daytime   13/udp
netstat   15/tcp
qotd      17/tcp     quote
msp       18/tcp     # message send protocol
msp       18/udp
chargen   19/tcp     ttytst source
chargen   19/udp     ttytst source
ftp-data  20/tcp
ftp       21/tcp
fsp       21/udp     fspd
ssh       22/tcp     # SSH Remote Login Protocol
ssh       22/udp
telnet    23/tcp
smtp      25/tcp     mail
time      37/tcp     timserver
time      37/udp     timserver
rtp       39/udp     resource# resource location
nameserver 42/tcp     name      # IEN 116
whois     43/tcp     nickname
tacacs    49/tcp     # Login Host Protocol (TACACS)

```

Another important file is `/etc/protocols` (you may need to install *nmap* for an extensive list). You can use these files to translate protocol names to numbers (so that the IP layer on other hosts can understand):

```

ip 0 IP # internet protocol, pseudo protocol number
#hopopt 0 HOPOPT # IPv6 Hop-by-Hop Option [RFC1883]
icmp 1 ICMP # internet control message protocol
igmp 2 IGMP # Internet Group Management
ggp 3 GGP # gateway-gateway protocol
ipencap 4 IP-ENCAP # IP encapsulated in IP (officially "IP")
-----[output truncated]-----

```

The `getprotobyname()` function is the one that gives a `protoent` structure for the line from `/etc/protocols` by matching it with a protocol name. We can show the function as follows:

```

#include <netdb.h>

struct protoent *getprotoent(void);
struct protoent *getprotobyname(const char *name);
struct protoent *getprotobynumber(int proto);
void setprotoent(int stayopen);
void endprotoent(void);

```

The system uses `/etc/nsswitch.conf` to configure which services are to be employed by it in order to find information such as host names, password files, and

group files. Here is a typical entry:

```

passwd:    compat
group:     compat
shadow:    compat

hosts:     files mdns4_minimal [NOTFOUND=return] dns mdns4
networks:  files

protocols: db files
services:  db files
ethers:    db files
rpc:       db files

netgroup:  nis

```

We can use the *ethtool* command followed by the interface name to display the Ethernet card settings:

```

aasisvinayak@GNU-BOX:~$ sudo ethtool eth0
Settings for eth0:


Supported ports: [ TP MII ]
Supported link modes:  10baseT/Half 10baseT/Full
                      100baseT/Half 100baseT/Full

Supports auto-negotiation: Yes
Advertised link modes: 10baseT/Half 10baseT/Full
                      100baseT/Half 100baseT/Full

Advertised auto-negotiation: Yes
Speed: 100Mb/s
Duplex: Full
Port: MII
PHYAD: 0
Transceiver: internal
Auto-negotiation: on
Supports Wake-on: pumbg
Wake-on: g
Current message level: 0x00000033 (51)
Link detected: yes

```

This tool has a wide range of applications—check the man page for details.

We have covered some of the fundamentals in networking, though we still have a lot more basic stuff to discuss—say files like `/etc/inetd.conf`, `/etc/securetty`, tools like *tcpd* (access control facility), configuration of `/etc/hosts` and so on. We will cover these in the next edition. Then we shall move on to kernel-specific zones. Happy kernel hacking! 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

Post: Linux Administrator

Company: LIVIA India Pvt. Ltd.
Profile: Looking for candidates with an experience strictly in linux administration.
Exp: 4-7, **Location:** Mumbai
Email: Careers@livialegal.com

Post: Sr. Programmer Analyst

Company: Stream Net (India) Pvt. Ltd.
Profile: Required candidate shall be BE/ B.Tech/ equivalent degree from a reputed university.
Exp: 4-7, **Location:** Bengaluru
Email: jobs@streamnet.com

Post: Linux / Solaris / Virtualization Presales Consultant

Company: Oracle India Pvt. Ltd.
Profile: Technical background in Oracle open source products (Linux, Open Solaris, Oracle VM i.e. virtualization and database) focusing into Open source strategies. Good understanding of Linux / Solaris administration is desired with good exposure into Virtualization & VMWare/Redhat certification will be an added advantage.
Exp: 3-7, **Location:** Delhi
Email: pranesh.george@oracle.com

Post: Sr. Developers (C/ C++/ Linux)

Company: PT Pure Testing Software Pvt Ltd
Profile: Should have 2-4 yrs of experience in development having excellent programming skills in C or C++. Should be good expertise on Linux.
Exp: 2-4, **Location:** Noida
Email: amit@puretesting.com

Post: System Administrator- Web application

Company: AOL Online India Pvt Ltd
Profile: Strong experience on Linux based systems & Strong experience on Apache Tomcat is desired, alongwith Knowledge on MySQL administration (Basic) & Scripting skills (PERL / Shell) is desired.
Exp: 4-8
Location: Bengaluru
Email: pankaja.sridhar@corp.aol.com

Post: Java Developer

Company: Knowledge Quest Infotech Pvt. Ltd
Profile: Should be able to write functional specs, come up with high level and detailed design document.
Exp: 2-6
Location: Pune
Email: naukri@kqinfotech.com

Post: Sr. Network Administrator

Company: FCS Software Solutions Ltd.
Profile: Sound knowledge of Routing, switching and security is a must. Hands on experience on IPS/IDS, ASA, Firewalls, higher end L3 switches, routers is also desired.
Exp: 5-7, **Location:** Pune
Email: nmaheshwari@fcsltd.com

Post: Linux Administrator

Company: New Age Sys Solutions Pvt Ltd.,
Profile: Looking for Red Hat certified people having an experience with web hosting & web development platform and undertake system admin work in Linux/Unix Installation and configuration of Linux server and Linux client machine.
Exp: 2-5, **Location:** Ernakulam
Email: resumes@newagesysindia.com

Post: Unix/ Linux Administrator

Company: TNS India Ltd.
Profile: Required candidate shall be any Graduate with an experience of 3+ years in Network Centre Operations (NOC). Good working experience on Unix (preferably HP- UX), Linux , Solaris Administration are essential. Should have basic knowledge in Netback up process on operating systems.
Exp: 3-6, **Location:** Hyderabad
Email: orsc.vacancy@tns-global.com

Post: System Engineer(L1)- Linux

Company: Estel Technologies Pvt. Ltd.
Profile: Qualifications required:- BE/B.Tech/MCA/Graduate or equivalent & should be Redhat certified. 1-2 yrs experience of Linux Administration or Application support is desired. Should have good knowledge in Linux with knowledge on scripting.
Exp: 1-2, **Location:** Chennai
Email: jobs@esteltelecom.com

Post: Technical Support- Linux Professional

Company: Libsys Technologies Pvt. Ltd.
Profile: Will be required to provide technical support on LibSys software on Windows/ Linux/ Unix to sites in their region through email/ phone/ visits, do installation at new sites, Provide briefing, Maintaining healthy rapport with the Clients etc.
Exp: 2-7, **Location:** Gurgaon
Email: hrd@libsys.co.in

Post: Sr. Infrastructure Specialist- Linux

Company: Yash Technologies Pvt Ltd
Profile: Experience in administrating RedHat Linux on various hardware- HP, IBM, etc. Working experience on fast phased 24X7 production support is desired.
Exp: 5-10
Location: Pune
Email: anil.kola@yash.com

Post: Linux System Administrator

Company: Flagstone Underwriting Support Services (India) Pvt Ltd
Profile: 5 years of experience in managing RedHat / CentOS Linux servers in a development and production environment is desired. Excellent familiarity with RedHat / CentOS based Linux distributions & shell scripting & installing non-standard software from source code & familiarity with standard server software (Apache, BIND, Postfix, Sendmail, yum, LDAP, NFS, MySQL, PostgreSQL) is desired.
Exp: 5-8
Location: Hyderabad
Email: india-resumes@flagstonere.bm

Post: System Administrator- Linux

Company: Pramaya Knowledge Solutions Pvt Ltd
Profile: Following skills are desired: Linux (RHEL, CentOS, Fedora, Ubuntu) alongwith Basic Networking skills- MS Windows XP/Vista. Must have one or more of Red Hat Linux Based Certifications.
Exp: 2-5
Location: Bengaluru
Email: careers@pramata.com



Who says sky is the limit?

Explore endless career opportunities on Naukri.com

FOSS Yellow Pages


The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

✂ Tear & Send

ORDER FORM

Tear & Send ✂

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB-HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of

EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. _____

against my credit card No. _____

C V V No. _____ (Mandatory)

Date of Birth ____ / ____ / ____ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)

Signature (as on the card)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Consultant (Firm)

IB Services

Free Installation of GNU/Linux on Laptops and Desktops. Thin client solutions based on Debian and Ubuntu. Laptops and Desktops pre-installed with Debian and Ubuntu. Migration to GNU/Linux. Data Recovery. Navi Mumbai
Kerala
Mobile: 09847446918
Email: ibmanoj@gmail.com
Web: www.ibservices.in

OS3 Infotech

- Silver Solutions Partner for Novell
- High Availability Computing Solutions
- End-to-end Open Source Solutions Provider
- Certified Red Hat Training Partner
- Corporate and Institutional Training

Navi Mumbai
Mobile: 09324113579
Email: info@os3infotech.com
Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX, OSS & JBOSS related projects. We specialize in high-availability and high-performance clusters, remote and onsite system management, maintenance services, systems planning, Linux & JBOSS consulting & Support services.
Hyderabad
Mobile: 09392493753, Fax: 040-40131726
Email: manojkummar@taashee.com
Web: www.taashee.com

Torrid Networks Pvt Ltd

Torrid is a leading provider of information technology consulting with focus on Information Security Services and Open Source Solutions.
Noida
Tel: 0120-4545100, Mobile: 09015505583
Fax: 0120-4235064
Email: tarun.kassana@torridnetworks.com
Web: www.torridnetworks.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek vibrain B1 UMPC with Linux, Touch Screen, 1 GB RAM 60GB, Wi-Fi, Webcam, upto 6 hour battery (opt.), USB Port, max 1600x1200 resolution, screen 4.8", 7.5"x3.25" Size, weight 526 gm.
New Delhi

Mobile: 09968756177, Fax: 011-26187551
Email: comptekdelhi@compteki.com
Web: www.compteki.com or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation Education and Training
Mumbai
Tel: 022-28272300, 66462300
Fax: 022-28272399
Email: customer@aptech.ac.in
Web: www.aptech-education.com, www.arena-multimedia.com

IT-Campus: Academy of Information Technology

IT training and solution company with over 12 years of experience. - RHCE
• Software Training • Hardware Training
• Multimedia And Animation • Web Designing • Financial Accounting
Kota (Raj.)
Tel: 0744-2503155, Mobile: 09828503155
Fax: 0744-2505105
Email: m_trilok@yahoo.com
Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization is primarily engaged in Education and Training through its own & Franchise centres in the areas of IT Software, Hardware, Networking, Retail Management and English. The institute also provides customized training for corporates.
New Delhi
Tel: 011-25916832-33
Email: info@mahanindia.com
Web: www.mahanindia.com

Enterprise Comm. Solutions

Aware Consultants

We specialize in building and managing Ubuntu/Debian Linux servers and provide good dependable system administration. We install and maintain in-house corporate servers. We also provide dedicated and shared hosting as well as reliable wireless/hybrid networking.
Bangalore
Tel: 080-26724324
Email: sales@aware.co.in
Web: www.aware.co.in

Cynapse India Private Limited

We are the creators of open source

product cyn.in. cyn.in is a web 2.0 group collaboration software created by Cynapse, that inter-connects your people with each other and their collective knowledge, seamlessly. It combines the capabilities of collaboration tools like wikis, blogs, file repositories, micro blogs, discussions, audio, videos, and other social applications into a seamless platform. cyn.in helps teams to build collaborative knowledge by sharing and discussing various forms of digital content within a secure, unified application that is accessible using a web based interface or a rich desktop client.
Mumbai
Tel: 022-28445858, 28445629
Email: lisa@cynapse.com
Web: www.cynapse.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.
Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty at the Indian Institute of Science, Bangalore and carry over eight decades of experience and fundamental knowledge in the field of DSP and Telecommunication. ESQUBE plays a dominant role in the creation of IP in the domain of Sensors, Signals and Systems.
Bangalore
Tel: 080-23517063
Email: info@esqube.com
Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open source Enterprise Groupware
Mobile push, Email Syncing of Contacts/Calendar/Tasks with mobiles
• Mail Archival • Mail Auditing • Instant Messaging
New Delhi
Tel: 011-30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

netCORE
The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux Based Enterprise Mailing solution with around 1500+ Customer all over India. Key Solutions:
• Enterprise Mailing and Collaboration Solution • Hosted Email Security • Mail Archiving Solution • Push Mail on Mobile • Clustering Solution
Mumbai
Tel: 022-66628000
Mobile: 09322985222
Email: kalpit@netcore.co.in
Web: www.netcore.co.in

 **redhat.**

Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.
Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis TechnoLab is the first of its kind, state-of-the-art infrastructure, Hardware, Networking and I.T Security training institution headquartered in Kolkata. TechnoLab is the training division of Xenitis group of Companies. It is the proud owner of 'Aamar PC', the most popular Desktop brand of Eastern India. These ranges of PC's are sold in the west under the brand name of 'Aamchi PC', in the north as 'Aapna PC' and in the south as 'Namma PC'.
Kolkata
Tel: 033-22893280
Email: srinku@xenitisgroup.com
Web: www.technolabindia.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873839960

Email: sales@aisplglobal.com

Web: www.aisplglobal.com

Advent Infotech Pvt Ltd

Advent has an experienced technomarketing team with several years of experience in Networking & Telecom business, and is already making difference in market place. ADVENT qualifies more as Value Added Networking Solution Company, we offers much to customers than just Routers, Switches, VOIP, Network Management Software, Wireless Solutions, Media Conversion, etc.

New Delhi

Tel: 46760000, 09311166412

Fax: 011-46760050

Email: marketingsupport@adventelectronics.com

Web: www.adventelectronics.com

Asset Infotech Ltd

We are an IT solution and training company with an experience of 14 years, we are ISO 9001: 2000. We are partners for RedHat, Microsoft, Oracle and all Major software companies. We expertise in legal software ans solutions.

Dehradun

Tel: 0135-2715965, Mobile: 09412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in

Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veritas, Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

HBS System Pvt Ltd

System Integrators & Service Provider.Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symantic Partners of SUN for their new startup E-commerce initiative Solution Provider on REDHAT, SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amittal@hbsindia.com

IgnitionWorks

IgnitionWorks provides customized IT solutions for small/medium

businesses by leveraging open source technologies with excellent support & affordable costs. Our expertise include custom software development & infrastructure deployment/support for all industries.

Bangalore

Tel: 080-41243227, Mobile: 9611121452

Email: info@ignitionworks.net

Web: www.ignitionworks.net

INGRES®

Business Open Source

Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

Is the authorised partner for RedHat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Pacer Automation Pvt Ltd

Pacer is leading providers of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, Vmware, Microsoft and Symantec. Our core expertise exists in, Consulting, building and Maintaining the Complete IT Infrastructure.

Bangalore

Tel: 080-42823000, Fax: 080-42823003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

Srijan Technologies Pvt Ltd

Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

industry verticles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tetra.in

Web: www.tetra.in

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

ManasHosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd

OpenLx Linux with Kalcutate (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

Linux Experts

Intaglio Solutions

We are the training and testing partners of RedHat and the first to conduct RHCSS exam in delhi for the first time ever.

New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845009939, +91-9343861758

Email : sales@gte-india.com

Web: www.gte-india.com

Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/ DVDs with each purchase. Visit our website for product and membership details

Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in



InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

Noida

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: http://opensource.infoaxon.com

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Sarvasv Technologies Pvt Ltd

Web OSX Windows Linux Mobile LAMP W3C apps. Agile, BDD, TDD development. On-Demand, SaaS, ERP, CRM, HRM. Open Source integration. Online project management. VCS (git, svn). automated deployment.

Greater Noida

Mobile: 09711472457, 09810270848

Email: sales@sarvasv.in

Website: www.sarvasv.in

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patiala

Mobile: 91-9780531682

Email: pardeep@mrsite.co.in

Web: www.mrsite.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648668, 66091565

Email: ceo@salahsoftware.com

Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing

Mobile: 09880715253

Email: ranga@digitalwaves.in

Web: www.digitalwaves.in

Enjoy Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India.

Gujarat

Tel: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjoyworld.com

Web: www.enjoyworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also

providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455-66

Fax: 079-27545488

Email: info@bascombridge.com

Web: www.bascombridge.com

Brainnet

Kolkata

Tel: 033-40076450

Email: brainnet@brainware-india.com

Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586

Fax: 022-25706700

Email: contact@cettm.mtnl.in

Web: http://cettm.mtnl.in/infra

Complete Open Source Solutions

RHCT, RHCE and RHCSS training.

Hyderabad

Tel: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com

ElectroMech

Redhat Linux and open source solution, RHCE, RHCSS training and exam center, Ahmedabad and Vadodara

Ahmedabad

Tel: 079-40027898

Email: electromech@electromech.info

Web: www.electromech.info

Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2335324

Email: enquiry@focuzinfotech.com

Web: www.focuzinfotech.com

G-TEC Computer Education

ISO 9001:2000 certified IT Company, International Testing Centre, Specialised in Multimedia & Animation, conduct MCP, MCSE 2000, MCDBA and MCSA certificates, CCNA, CCNP, the Only authorized centre by INTERNATIONAL AND EUROPEAN COMPUTER UNION to conduct ICDL, Adobe Certifications, training on Web Designing, Tally, Spoken English. Conducts Corporate and institutional training. International certifications issued.

Bangalore

Tel: 080-43567000

Email: gtec.indiranagar@gmail.com

Gujarat Infotech Ltd

GIL is a IT compnay and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubantoo, and PHP, Mysql

Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: info@gujaratinfotech.com

Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

Chennai

Tel: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking, while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525

Email: info@mazenetsolution.com

Web: www.mazenetsolution.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning. Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing Corporate and Open classes for RHCE / RHCSS training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

New Delhi

Tel: 46526980-2

Mobile: 09310024503, 09312411592

Email: info@networknuts.net

Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile:

09873108801

Email: rakhi@stg.in

Web: www.stgonline.com

www.stgglobal.com

TNS Institute of Information Technology Pvt Ltd

Join RedHat training and get 100% job guarantee. World's most respected Linux certification. After RedHat training, you are ready to join as a Linux Administrator or Network Engineer.

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

New Delhi

Tel: 011-3085100, Fax: 30851103

Email: nks@tiit.co.in

Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

Kolkata

Tel: 033-22833568, Mobile: 09433111110

Email: enquiry@webelinformatics.com

Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Redhat Training Partner in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses—RHCSS, SCNA, MCSE, CCNA.

Kolkata

Tel: 033-25488736, Mobile: 09830075018

Email: sinhatuhin1@gmail.com

Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09760099050

Web: www.agamtecindia.com

Amritha Institute of Computer Technology

Amritha Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.

Mobile: 09393733174

Email: aict.hybd@amrita.ac.in

Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to

the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.

Bhopal

Tel: 0755-2661412, 2661559

Fax: 0755-4220022

Email: crisp@crispindia.com

Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09312506496

Email: info@cfosdr.com

Web: www.cfosdr.com

Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSS, CCNA, MCSE

Kolkata

Tel: 033-25395508, Mobile: 09831705913

Email: info@cisconetinfo.com

Web: www.cisconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available..Leading Networking Institute in Marathwada

Aurangabad

Tel: 0240-3299509, 6621775

Email: aurangabad@cmsinstitute.co.in

Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

Bikaner

Tel: 0151-2202105, Mobile: 09928173269

Email: cmtech.bikaner@gmail.com,

kr.gupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer

Hardware & Networking

Dehradun

Tel: 3208054, 09897168902

Email: thedishainstitute@gmail.com

Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux-Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067566-67, 2609849

Fax: 0172-2615465

Email: info@technoschool.net

Web: http://technoschool.net

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications

Jaipur

Tel: 0141-3213378

Email: franchise_gt@gteducation.net

Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

Bhopal

Tel: 0755-4094852

Email: bhopal@hclcdc.in

Web: www.hclcdc.in

IINZTRIX E Technologies Pvt Ltd

No. 1 Training prvider in this region.

meerut

Tel: 0121-4020111, 4020222

Mobile: 09927666664

Email: jai@iintrix.com

Web: www.iintrix.com

Indian Institute of Job Oriented Training Centre

Ahmedabad

Tel: 079-40072244—2255—2266

Mobile: 0989749595

Email: info@iijt.net

Web: www.iijt.net

Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking
•Software Engg. •Multimedia Training.

Ahmedabad

Tel: 079-32516577, 26607739

Fax: 079-26607739

Email: contact@iantindia.com

Web: www.iantindia.com

IPCC

Bridging Gap with professionals.

Lucknow

Tel: 0522-3919496

Email: ipccsko@yahoo.co.in

Web: www.ipcc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.

New Delhi

Mobile: 09910710143, Fax: 011-25886909

Email: info@koenig-solutions.com

Web: www.koenig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professional & Cooperate.

Meerut

Tel: 0121-2420587, Mobile: 9997526668

Email: Info@nacsglobal.com

Web: www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is an organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.

Meerut

Tel: 0121-2767756, Fax: 0121-4006551

Mobile: 09897796603

Email: info@nacsglobal.com,

mohit@nacsglobal.com.

Web: www.nacsglobal.com

Netdiox Computing Systems

We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.

Bangalore

Tel: 080-26640708

Mobile: 09740846885

Email: info@netdiox.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

NetMax-Technologies

Training Partner of RedHat, Cisco
Chandigarh
Tel: 0172-2608351, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com

Netxprt institute of Advance Networking

Netxprt Noida is a Leading organization to provide Open Source training on RedHat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.

Noida
Tel: 0120-4346847, Mobile: 09268829812
Email: anoop.vyas@netxpriindia.com
Web: www.netxpriindia.com

Netzone Infotech Services Pvt Ltd

Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.

New Delhi
Tel: 011-46015674, Mobile: 9212114211
Email: info@netzoneindia.net

Neuron IT Solutions

We offer end to end services and support to implement and manage your IT Infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration.

Chennai
Mobile: 09790964948
Email: enquiry@neuronit.in
Web: www.neuronit.in

Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the

unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.
Chennai
Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education

RHCE & RHCSS Certifications
Jabalpur
Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.

Bangalore
Tel: 080-26639207, 26544135, 22440507
Mobile: +91 9945 282834
Email: counsellors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.

Chandigarh
Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre

SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT)

from KARNATAKA STATE OPEN UNIVERSITY.
New Delhi
Tel: 011-26183327, Fax: 011-26143642
Email: southdelhicomputercentre@gmail.com, southdelhicomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Ssytms Quest

Making Tomorrow's professionals TODAY
Bangalore
Tel: 080-41301814
Email: directorv@ssystemsquest.com
Web: www.ssytmsquest.com

Trimax FuturePerfect

A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.
Mumbai
Tel: 022-40681313, Mobile: 09987705638
Fax: 022-40681001
Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award " REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.
Mumbai
Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses adn authorised testing center of Prometric, Vue and Red Hat.
Mumbai
Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com

To advertise in this section,
please contact
Somaiah (B'lore):
09986075717) Dhiraj
(Delhi: 09811206582) on
011-2681-0602 Extn. 222

Yash Infotech

Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.
Lucknow
Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Web Hosting

IDS Logic Pvt Ltd

IDS Logic is a leading innovative IT Solutions company provides Linux Web Hosting, Windows Web Hosting, Application Hosting, Reseller Hosting, Dedicated Server, Virtual Private Server, Server Management & Monitoring, Domain Name Registration, SSL Certificate, Corporate Mailing Solutions, Web Security Solutions.
Noida
Tel: 0120-4235665, Fax: 0120-4235665
Email: info@idslogic.com
Web: www.idswebhosting.com

Perfect Innovation

Web Hosting Spider
• Web Hosting • Web Design
• Web Application Development
• SMS Hosting • Corporate Hosting
• Dedicated Servers
Puducherry
Tel: 0413-3202726, 3246999
Email: admin@webhostingspider.com
Web: www.webhostingspider.com

Want to register your organisation in
FOSS Yellow Pages For **FREE***

Call: Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

*Offer for limited period.

new year

Subscription Offer

Subscribe and Select Your Gift!



Magazine	Cost of 12 issues (Rs)	You Save (Rs)	You Pay (Rs)	SELECT YOUR GIFT
Electronics For You (with CD) <input type="checkbox"/>	720	145	575	Please mark three of the below as your 1st, 2nd and 3rd choice, in boxes. <input type="checkbox"/> Learn Macromedia Flash Mix CD <input type="checkbox"/> Tutor C# CD <input type="checkbox"/> Tutor Net CD <input type="checkbox"/> Wrist Watch, EFY <input type="checkbox"/> Electronics Projects Volume Handbook (Please specify any Volume No. from 1 to 24) <input type="checkbox"/> ABC of Amateur Radio Book
Electronics Bazaar <input type="checkbox"/>	600	120	480	
Linux For You (with CD & DVD) <input type="checkbox"/>	1200	240	960	
BenefIT <input type="checkbox"/>	600	125	475	
Facts For You <input type="checkbox"/>	1200	250	950	

hurry! The Gifts Stocks are limited

Gifts will be sent on a first come first served basis, till stocks last. You are likely to receive some other gift (such as pen set or 'T' shirt) when stocks for the above finish, in case we get your order with payment by 31st January, 2010

☒ **Don't Miss A Single Issue!**

☒ **Get Your Money Back if Not Satisfied****

☒ **Ensure Regular Supply**

☒ **Order Now For FREE Home Delivery**

☒ **FREE Replacement Policy***

* Replacement will be made upon receipt of the damaged / non-receipt copies within 30 days of its publication
 ** After three months, if you are not satisfied with the magazine, your balance amount will be returned.



Name _____ Company Name _____ Designation _____

Mailing Address _____ Distt. _____

City _____ Pin Code _____ State _____ Phone _____

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓)marked magazine(s), for the period _____. Please find enclosed a sum of Rs _____ by DD/MO/crossed cheque* bearing the No. _____ dt. _____ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs 50 on non-metro cheque)**

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. _____

against my credit card No. _____

Date of Birth ____/____/____ (dd/mm/yy) Card Expiry Date ____/____ (mm/yyyy)

Signature (as on the card) _____


Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers ordering through this form. # The rates are valid for subscribers within India only. # Please allow 4-6 weeks for processing of your subscription and the gift. # Gift against cheque payment will be sent after clearances of the cheque # No discounts will be available for subscription agencies under this scheme. # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.



Send this filled-in form or its photocopy to :

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020; Ph: 011-26810601-03; Fax: 011-26817563; e-mail: info@efyindia.com; website: www.efyindia.com

VALID TILL JANUARY 31, 2010 ONLY



EnterpriseDB[®]

The Enterprise Postgres Company

in association with



Future is open

Lynus Academy

The official EnterpriseDB PostgreSQL Training

Topics covered:

- Introduction and Architecture ● Installation ● Introduction to SQL
- Performance Tuning ● Routine Database Maintenance Tasks
- Backup and Recovery ● High Availability & Replication
- Security ● PostgreSQL Partitioning and Tablespaces
- Monitoring and Statistics ● Database Migration

Dates

18th, 19th and 20th
of December, 2009

From 10 AM to 5 PM

Location

1961B, Vijaya Complex 5th Floor
(Near - Thirumangalam Signal &
Indian Bank), Asiad Colony,
Anna Nagar - West,
Chennai - 600 101, Tamilnadu, India

To register e-mail us at

training@enterprisedb.com
contactus@lynusacademy.com

Call us on

Pune - 020-30589500/ 01
Chennai - 044- 42171278



Future is open

Lynus Academy

EnterpriseDB[®]

The Enterprise Postgres Company

www.enterprisedb.com



MOBLIN™ ZONE



MOBLIN™ ZONE

Create rich "PC-like" internet and media experience on mobile devices with Moblin™ v.2

Moblin™ v.2 is an optimized open source Linux software stack and application framework for enabling rich Internet experiences on Intel® Atom™ Processor based MID, netbooks, nettops, IVO, and embedded systems.

Moblin™ v.2 provides the most flexible application framework including a 3D UI for developing Rich Internet and media applications that run across multiple Intel® Atom™ Processor based devices with minimal UI adaptation. The new Moblin™ v.2 SDK includes reference code, documentation, and tools to simplify development and provide faster TTM.

To learn more about Moblin™ v.2, visit www.moblinzone.com and register to receive regular updates.

Sell your Moblin™ Apps through the Intel® Atom™ Developer Program.
Earn up to 70% revenue on every sale.
Learn more at appdeveloper.intel.com